American Economic Journal: Economic Policy 2020, 12(3): 279–311
https://doi.org/10.1257/pol.20180198

Price-Linked Subsidies and Imperfect Competition in Health Insurance†

By Sonia Jaffe and Mark Shepard*

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 postsubsidy 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 percent, and much more in less competitive markets. For cost uncertainty reasonable in a mature market, these losses outweigh the benefits of price-linking. (*JEL* G22, H75, I13, I18)

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

* Jaffe: Microsoft, One Microsoft Way, Redmond WA 98052 (email: spj@uchicago.edu); Shepard: John F. Kennedy School of Government, Harvard University, 79 JFK St, Mailbox 114, Cambridge MA 02138, and NBER (email: mark_shepard@hks.harvard.edu). Matthew Notowidigo was coeditor for this article. 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 Northwestern University 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. Jaffe gratefully acknowledges the support of a National Science Foundation Graduate Research Fellowship. 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. All mistakes are our own. Replication code and the publicly available portion of our dataset are available at https://www.openicpsr.org/openicpsr/workspace?goToPath=/openicpsr/115405.

† Go to https://doi.org/10.1257/pol.20180198 to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

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. This choice between subsidy types is also relevant to some noninsurance markets where the government subsidizes individuals' purchases from private providers, such as school vouchers and housing subsidies.

While linking subsidies to prices can be convenient for regulators, 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 model to formalize the competitive implications of price-linked subsidies. We do so in a setting, modeled on the ACA case, where all competing firms receive the *same subsidy amount*—a key principle of "managed competition" (Enthoven 1988) intended to avoid the clear distortions when each firm's price determines its *own* subsidy (e.g., percentage subsidies).[1] We show that market-level price-linked subsidies still weaken price competition when there is 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.[2]

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 (i) 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) or (ii) has more market power (a smaller own-price elasticity of demand). The outside option of uninsurance is important in the ACA case since a large share of the subsidy-eligible population remains uninsured. Furthermore, the ACA exchanges are highly concentrated, with about half of consumers living in markets with just one or two insurers as of 2018 (Kaiser Family Foundation 2018).

[1] Percentage-of-price subsidies are commonly used in employer health insurance and have been shown to weaken price competition (Cutler and Reber 1998, Liu and Jin 2015) though also to mitigate problems with adverse selection.

[2] While fixed subsidies can affect prices by shifting the demand curve (Decarolis, Polyakova, and Ryan forthcoming), they do not change the *slope* of the demand curve in this way that clearly distorts prices upward.

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, 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, and a pure gain for consumers. 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 V. We also pay special attention to 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. We propose (and include in our model) two reasons why optimal subsidies may vary with prices. First, the government may wish to insure low-income enrollees against the risk of price shocks—the ACA's subsidies were intended to ensure that postsubsidy premiums are "affordable" regardless of insurance prices. Second, higher prices may reflect higher underlying health care costs and therefore higher costs of uncompensated care for the uninsured—an externality borne by hospitals and clinics that the uninsured do not internalize when choosing to forgo coverage.[3] Price-linking automatically connects subsidies to market-level health care cost shocks if insurer prices reflect information on costs beyond what regulators have (or can use).

To study empirically the pricing distortions and welfare tradeoffs of price-linked subsidies, 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, low-income population, and price-linked subsidy design. We use our model of insurer competition to estimate the effects of price-linked subsidies based on parameters estimated from the Massachusetts data.

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 percent. We also use a difference-in-difference approach based on within-plan differential price changes to estimate an own-price semi-elasticity of demand of -2.16 percent.

[3] There is growing evidence on the importance of uncompensated care for low-income people (Finkelstein, Hendren, and Luttmer 2015; Garthwaite, Gross, and Notowidigdo 2018). Mahoney (2015) proposes including these costs (which he connects with the threat of bankruptcy) as a rationale for the mandate penalty.

Our second method is to use structural demand and cost models to simulate equilibrium 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 approach is that we use estimates from the natural experiments to identify key model parameters that govern substitution patterns among plans and between plans and uninsurance.

Across both methods, we find three sets of results. First, we estimate that price-linked subsidies raise the price of the pivotal plan by a nontrivial amount—by 9 percent (or \$36 per month) in the sufficient statistics method and 1–6 percent (or \$4–26 per month) in the simulations. Based on our 2011 simulations, switching from price-linked to fixed subsidies could achieve either the same insured rate at 6.1 percent lower subsidy cost or 3.1 percentage points greater insurance coverage at the same cost. Of course, the actual distortion for the ACA might be higher or lower because of the ways the ACA differs from Massachusetts.

Second, we show that the pricing distortion depends critically on the intensity of market competition. We find that the effect of price linking is about twice as large (6–12 percent 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. These two-competitor simulations also show that with four major insurers, switching to price-linked subsidies has a comparable effect on the price of the cheapest plan as removing all but one of that 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 uncertainty must be quite high—the regulator's cost prediction error must exceed $+/-12.5$ percent—for the benefits of price-linked subsidies to outweigh the losses from higher prices. This degree of uncertainty is unlikely in markets where lagged cost data are available, though it is more plausible for a new market.

This analysis casts doubt on cost uncertainty as a sufficient explanation for price-linked subsidies. We conclude that rationalizing price linking would require political economy factors. These factors could include limits on regulators' bandwidth to optimize fixed subsidies based on local health care costs, an extreme concern that postsubsidy 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 postsubsidy premiums and mitigated what might have been an adverse selection death spiral.

Related Literature.—Our paper is related to a small but growing body of research studying the (often unintended) competitive implications of subsidy policies—including in Medicare Part D (Decarolis 2015; Decarolis, Polyakova, and Ryan forthcoming), Medicare Advantage (Curto et al. 2014), and employer insurance (Cutler and Reber 1998, Liu and Jin 2015). Most closely related is concurrent work by Tebaldi (2017), which 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, including in Massachusetts (Ericson and Starc 2015, 2016). Recent work by Shepard (2016) and Finkelstein, Hendren, and Shepard (2019) uses the same CommCare setting and data as us, but it focuses on different research questions: Shepard (2016) studies hospital networks and adverse selection, while Finkelstein, Hendren, and Shepard (2019) quantifies willingness to pay for insurance.

The remainder of the paper is structured as follows. Section I derives the distortion effect and shows the welfare tradeoffs of price-linked subsidies in a simple model. Section II describes our setting and data, and Section III uses estimates from natural experiments to approximate this pricing distortion. Section IV estimates the structural model and presents simulations. Finally, Section V discusses the role of uncertainty and political constraints as well as the applicability of our framework to other markets, and Section VI concludes.

I. 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 online Appendix A1, we show that the basic logic of the distortion carries through to a more general model allowing for multiplan insurers (as in the ACA).

Insurers $j = 1, \ldots, J$ each offer a differentiated plan and compete by setting prices $P = \{P_j\}_{j=1,\ldots,J}$. The exchange collects these price bids and uses a prespecified formula to determine a subsidy $S(P)$ that applies equally to all plans.[4] Subsidy-eligible consumers then choose which (if any) plan to purchase based on plan attributes and postsubsidy 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, $Q_j\left(P^{cons}, M\right)$, is a function of all postsubsidy premiums and the mandate penalty.

We assume that insurers set prices simultaneously to maximize static profits, knowing the effect of prices on demand and cost. As in most recent work on insurance (e.g., Handel, Hendel, and Whinston 2015), we assume fixed plan attributes

[4] The 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.

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.[5] The insurer profit function is

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

A necessary condition for Nash equilibrium is that each firm's first-order condition holds:[6]

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

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

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

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.[7] This assumption implies 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: $\sum_k \partial Q_j / \partial P_k^{cons} + \partial Q_j / \partial M = 0$, $\forall j$. Using this condition to simplify equation (2), we get

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

[5] In online Appendix A2, 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.

[6] This condition is still necessary for Nash equilibrium in a more complicated model where insurers also set quality characteristics. Thus, our main theoretical point about price-linked subsidies holds when quality is endogenous, though there may also be effects on quality and cost levels not captured in our model.

[7] This assumption is typically justified by the fact that prices are a small share of income. After subsidies, premiums in our setting are just 0–5 percent of income, and price differences are even smaller. 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 IIIA that demand responds similarly to an increase in the mandate penalty and a decrease in all premiums. 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.

The effective demand slope (for a firm's pricing equation) equals the standard slope of the demand curve, plus an adjustment if either S or M is linked to prices.

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.

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 $\partial Q_j/\partial M$.

A. *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, Polyakova, and Ryan (forthcoming)—but they do not *rotate* the demand curve. Markups are

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

where $\eta_j \equiv -(1/Q_j)(\partial Q_j/\partial P_j^{cons})$ 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 postsubsidy premium equals an (income-specific) "affordable amount," regardless of its presubsidy price, then

$$S(P) = \min_j P_j - AffAmt$$

so $\partial S(P)/\partial P_{\underline{j}} = 1$, 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:

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

where $\eta_{\underline{j},M} \equiv (1/Q_{\underline{j}})(\partial Q_{\underline{j}}/\partial M)$ is the semi-elasticity of demand for $\underline{j}$ with respect to the mandate penalty. The first-order condition for the nonpivotal 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 semielasticities 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:

$$(4) \qquad \Delta Mkup_{\underline{j}} = Mkup_{\underline{j}}^{PLink} - Mkup_{\underline{j}}^{F} = \frac{\eta_{\underline{j},M}}{\eta_{\underline{j}}\left(\eta_{\underline{j}} - \eta_{\underline{j},M}\right)} > 0,$$

$$(5) \qquad \%\Delta Mkup_{\underline{j}} = \frac{Mkup_{\underline{j}}^{PLink} - Mkup_{\underline{j}}^{F}}{Mkup_{\underline{j}}^{F}} = \frac{\eta_{\underline{j},M}}{\left(\eta_{\underline{j}} - \eta_{\underline{j},M}\right)} > 0,$$

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 semielasticities 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.[8]

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 affects the pivotal plan, though there may be strategic responses by other firms. In a model with uncertainty about others' prices (e.g., due to uncertainty about others' costs), the distortionary term $\eta_{\underline{j},M}$ would be weighted by the

[8] This 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.

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.

An Alternate Policy.—We also note an interesting implication of the price-linking term in equation (3). Fixed subsidies make this term zero—eliminating the competitive distortion—by making the S and M fixed with respect to prices. Another way to eliminate the distortion is to link subsidies to prices ($\partial S / \partial P_j > 0$) but link the mandate penalty to prices in the *opposite direction* ($\partial M / \partial P_j = -\partial S / \partial P_j < 0$). Intuitively, this eliminates the distortion because it means that insurer prices do not affect $S + M$, the *net* public incentive for consumers to buy insurance. We discuss this alternate policy in more detail in online Appendix A3.

B. *Welfare and Uncertainty*

One of our paper's objectives is to study the welfare implications of price-linked subsidies when regulators face uncertainty. We base our welfare analysis on a consumer surplus standard (common in antitrust analysis), adjusted to subtract net government spending (subsidy cost net of mandate revenue) and a social externality of uninsurance. Together, these form a regulatory objective that we refer to as "public surplus"; see online Appendix A4 for a formal specification. In our empirical work, we also consider an alternative where the regulator's objective includes insurer profits, as in a social surplus standard.

The public surplus objective captures two rationales for subsidizing insurance: adverse selection and negative externalities from individuals being uninsured. We allow for two types of externalities from 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 (Finkelstein, Hendren, and Luttmer 2015; Garthwaite, Gross, and Notowidigdo 2018; Mahoney 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.

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 and therefore premiums (since the subsidy is unchanged). This implies 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 those additional insurance purchases.

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 plan's optimal price is increasing in the marginal cost of the subsidy-pivotal plan.* This is related to prices being strategic complements. (It holds trivially with logit demand.)

(iii) *Welfare is decreasing in each plan's price.* This necessarily holds if the regulator puts sufficiently little weight on firm profits and can hold more generally.

See online Appendix A5 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, because of two potential advantages. First, price-linked subsidies can allow regulators to better calibrate the subsidy when (i) regulators are uncertain about market costs (or prices), *and* (ii) the *optimal* subsidies are higher in states of the world where prices are unexpectedly high. 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.

Second, with cost uncertainty, price-linked subsidies also have the benefit of stabilizing postsubsidy consumer prices, transferring the risk of cost shocks to the government; we think this is reflected in the political rhetoric around "affordability." We incorporate the value of this risk protection into our simulations; online Appendix A4 shows how.

II. Setting and Data

A. *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 percent 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 four to five 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-percent-of-poverty group that made all plans free, regardless of their presubsidy 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.[9] The price of uninsurance is the mandate penalty after its introduction in late 2007 (as discussed in Section IIIA). 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 postsubsidy premium of the cheapest plan).

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 percent). 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.

B. *Data and Samples for Structural Model*

Administrative data from the CommCare program[10] 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

[9] In 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 percent 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.

[10] These data were 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.

individuals, whether or not they chose to purchase insurance.[11] 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 US 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,[12] 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 online Appendix B1 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 percent of poverty) to simulate the 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 reenrolled 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 reweight 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 online Appendix B1 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 percent of poverty. We do so both because this matches the subsidy-eligible

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

[12] Because 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.

Table 1—Summary Statistics

	CommCare			ACS (uninsured)	
	All cost estimate	New enrollees demand estimate	Above poverty simulation	New enrollees demand estimate	Above poverty simulation
Counts					
Unique enrollees	455,556	326,033	253,200	—	—
Sample size	—	—	—	4,562	2,225
Average 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					
<100 percent poverty	48.8%	51.3%	0.0%	47.0%	0.0%
100 percent–200 percent	38.2%	35.2%	74.6%	29.4%	55.4%
200 percent–300 percent	13.0%	13.4%	25.4%	23.6%	44.6%
	CommCare			ACS (uninsured)	
Simulation sample (above poverty)					
Share		54.5%		45.5%	
Cheapest plan in-market share		43.9%			
Presubsidy 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	
Costs sample (all enrollees)					
Observed cost		\$372.81			
Predicted in average plan		\$373.24		\$366.77	

Notes: 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. "Presubsidy 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 that 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 that 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.

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 percent 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 percent in the full ACS data for Massachusetts. Of those who enroll, about 44 percent choose the cheapest plan; presubsidy monthly prices average \$399, but subsidies are quite large. Enrollees pay an average of \$46 per month; the average 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 IVB) for the uninsured are somewhat lower than observed (and predicted) costs for the insured, consistent with the uninsured being a slightly healthier population.

III. 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 online Appendix B2.

A. *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 percent 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 percent poverty) to \$52.50 (for 250–300 percent poverty). People earning less than 150 percent of poverty did not face a penalty.

There was a spike in new enrollees into CommCare for people above 150 percent 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

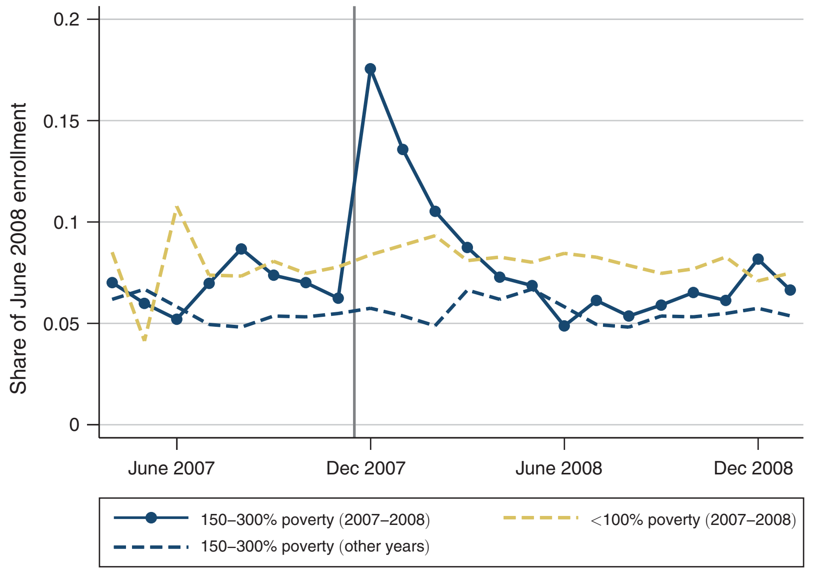

FIGURE 1. NEW ENROLLEES IN CHEAPEST PLAN AROUND MANDATE PENALTY INTRODUCTION

Notes: This graph shows monthly new enrollees (both first-time consumers and those reenrolling 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 December 2007 and the introduction of the mandate penalty, which applied only to the "150–300 percent poverty" income group (solid blue). The <100 percent poverty group (dashed yellow) is a control group not subject to the penalty. The "150–300 percent 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.

factors for this group at this time. As Figure 1 shows, there was no concurrent spike for people earning less than poverty (for whom penalties did not apply), and there was no enrollment spike for individuals above 150 percent of poverty in December–March of other years. Additionally, Chandra, Gruber, and McKnight (2011) shows 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 percent poverty and 150–300 percent poverty), we estimate

$$NewEnroll_{y,t} = \left(\alpha_{0,t} + \alpha_{1,t} \cdot Treat_y\right) \cdot MandIntro_t + \xi_y + \delta_y \cdot X_t + \zeta_{y,m(t)} \cdot DM_t + \varepsilon_{y,t},$$

where the dependent variable is new enrollment divided by that group's enrollment in June 2008, $MandIntro_t$ indicates months during the penalty introduction period (December 2007–March 2008), and $Treat_y$ is a dummy for the treatment group

(150–300 percent poverty). The regression also includes income group fixed effects (ξ_y), income group-specific year dummies and time polynomials ($\delta_y X_t$), and income-group specific dummies for December through March in all years ($\zeta_{y,m(t)} \cdot DM_t$)—which together make this a triple-difference specification.

The coefficients of interest are $\alpha_{1,t}$, which correspond to the excess enrollment for the treatment group during December 2007 to March 2008 in Figure 1. Column 3 of online Appendix Table B1 reports these estimates. Summing these $\alpha_{1,t}$ coefficients, we estimate that the mandate penalty caused a 22.2 percent increase in enrollment in the cheapest plan. As we describe in Section IV, we match the estimated 22.2 percent increase in enrollment as a moment in the structural model. Translating this increase into a semi-elasticity of demand, we find that enrollment in the cheapest plan increases by $\eta_{\underline{j},M} = 0.95\%$ on average for each \$1 increase in the penalty.

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 reestimate these regressions using enrollment in the *second-cheapest* CommCare plan as the outcome, we find very similar effects (an enrollment increase of 21.1 percent). 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, Gruber, and McKnight (2011) also studies the CommCare market and finds similar results, though it focuses on the effects of the mandate on adverse selection rather than the net increase in coverage. Hackmann, Kolstad, and Kowalski (2015) studies 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 percent 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 percent of poverty. Enrollees between 200–300 percent 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 online Appendix B2. 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 percent. In particular, the estimated semi-elasticity for the 150–200 percent 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.

B. *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 online Appendix B2. 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

$$(6) \qquad \log\left(NewEnroll_{j,y,r,t}\right) = \eta \cdot P^{cons}_{j,y,r,t,} + \xi_{j,r,t} + \xi_{j,r,y} + \epsilon_{j,y,r,t},$$

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, η corresponds to the semi-elasticity of demand with respect to own price (η_j in the theoretical model).

We get a semi-elasticity of 2.16 percent when weighted by average new enrollment in the region (see column 3 of online Appendix Table B3). We can compare these estimates to Chan and Gruber (2010), which studies price sensitivity of demand in this market using a different identification strategy. Though it reports an own-price semi-elasticity of demand of 1.54 percent, it does not allow for substitution to the outside option of uninsurance. When we adjust for that, its estimates imply a semi-elasticity for the cheapest plan of 2.17 percent.[13]

[13] The in-market share of the cheapest plan is 47.3 percent and 54.5 percent 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 $(1-0.473\times 0.545)/(1-0.473) = 1.41$, 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 $1.54\% \cdot 1.41 = 2.17\%$.

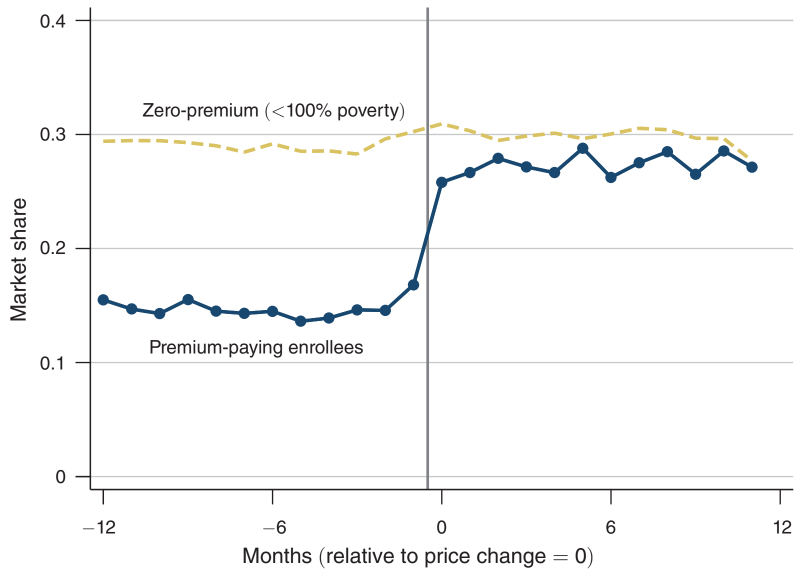

Figure 2. Identifying the Own-Price Semi-elasticity: Market Share around Price Decrease

Notes: 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 zero. The identification comes from comparing demand changes for above-poverty price-paying (new) enrollees (for whom premium changes at time zero) 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.

C. *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 IA to get an approximation of the price effect. The semi-elasticity estimates of $\eta_{j,M} = 0.95$ percent (Section IIIA) and $\eta_j = 2.16$ percent (Section IIIB) suggest that price-linked subsidies increase the price of the cheapest plan by \$36 (with a standard error of \$6.1)[14] per month (equation (4)). This is about a 9 percent 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 semielasticities are not constant, it will be less accurate for nonmarginal 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 one, the price distortion will be smaller. Second, this estimate does not allow for adverse selection into the

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

market, which would (i) reduce markups and therefore the dollar value of the main price distortion and (ii) 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 online Appendix A2 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.

IV. 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 II, and simulate equilibria under different subsidy policies and market structures.

A. *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 five-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.

The utility for consumer i of plan j equals

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

where $P_{j,i}^{cons}$ is the plan's postsubsidy 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 online Appendix C1. 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, Levinsohn, and Pakes (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 nonrandom 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, Levinsohn, and Pakes (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 IIIB.) 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 absorbs all price variation except for these differential changes across incomes, just as in a difference-in-difference model.

Two observations lend credence to this identification strategy. First, recall that a plan j's quality attributes (cost sharing, networks, customer service) are *identical* across incomes: different subsidies simply make the same plan have different premiums. Our setup allows for a region-year specific plan dummy, $\xi_{j,r,t}$, to capture this common quality. It also allows for persistent preference differences across income groups via the $\xi_{j,r,y}$ terms. It simply assumes that any *changes* in plan quality apply equally to all incomes within a region. Second, we can assess this common quality changes assumption by examining trends in plan shares before price changes—as in a difference-in-difference test for parallel pretrends. Figures 2 and B3 are strongly consistent with parallel (and basically zero) trends in market shares for the below- versus above-poverty groups prior to a price change.

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 IIIA. Specifically, we match the estimated 22.2 percent 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 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.

Demand Estimates.—We report full demand model estimates in online Appendix C2. The coefficients vary in sensible ways. Poorer consumers are more price sensitive, with baseline $\alpha(Z_i)$ declining by half from the 100–150 percent poverty to the 250–300 percent poverty group. Females and older consumers are less price sensitive, with the oldest group less than half as price sensitive as the youngest. The plan quality estimates (ξ_j) indicate that consumers see CeltiCare as inferior to other plans, with a quality shortfall worth about \$17–34 in monthly premiums (depending on consumers' $\alpha(Z_i)$). The other plans' average qualities are much more similar.

Table 2 shows the model's implied semielasticities with respect to own price and the mandate penalty. We find that each \$1/month increase in a plan's consumer premium lowers its demand by an average of 2.41 percent.[15] 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 margin $\left((P - mc)/P\right)$ of 11 percent above marginal cost, but adverse selection would imply a smaller margin with respect to average costs. Perhaps consistent with this, the average plan's actual margin above average cost in the data was \$399 – \$372 = \$27 (see Table 1), or 7 percent of revenue.

B. *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 IVC, risk adjustment works by adjusting plan revenues by multiplying price times a consumer risk score, which we model

[15] Our estimate is somewhat larger than the 1.5 percent estimate reported by Chan and Gruber (2010) for CommCare in an earlier period—even after adjusting its number to allow for substitution to uninsurance, which it does 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.

Table 2—Average Semielasticities

Own-price semi-elasticity				Semi-elasticity of insurance with respect to mandate penalty	
By plan (percent)		By year (percent)		By year (percent)	
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

Notes: 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.

separately. 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

$$E(c_{ijt}) = \exp(\mu X_{it} + \psi_{0,t} + \psi_{j,r,t}). \tag{7}$$

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 ($\psi_{j,r,t}$). The region-year plan effects capture cost implications of varying provider networks and 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 allows for adverse selection through the correlation between insurance demand and demographics.

A concern with a basic maximum likelihood estimation of equation (7) 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 use the panel data to estimate $\psi_{j,r,t}$ from *within-person* cost variation in a model with individual fixed effects (see online Appendix C1 for the method details). For this estimate, we limit the sample to individuals who are enrolled in one plan, then leave the market, and rejoin the market choosing a different plan (e.g., because plan prices have changed). While not perfect—there could be selection on risk changes over time—this method at least controls for unobserved risk differences that are stable over time.

Estimated Parameters.—Table 3 shows averages of plan cost effects, separately for before and after 2010 when CeltiCare entered. The numbers reported are percent effects on costs (i.e., $\exp(\psi_{j,r,t}) - 1$) and are relative to the share-weighted average plan in each year. Prior to 2010, both Network Health and BMC had similar cost

Table 3—Cost Parameter Estimates

	Plan effects				
	CeltiCare	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%

Notes: 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.

effects—about 7 percent below average, with other plans somewhat higher. When CeltiCare entered, it became the clear low-cost plan—32 percent below average. Cost effects of the other plans changed somewhat, but their ordering did not.

The other cost parameters are reported in online Appendix C2. 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.

C. *Simulation Methods*

We use these demand and cost models and estimates to simulate pricing equilibria under alternate subsidy policies. We compare the equilibrium under price-linked subsidies to fixed subsidies where each income group's subsidy equals the subsidy amount that emerged under the price-linked equilibrium. This ensures that we are comparing subsidy designs with equal-size subsidies but simply varying pricing incentives. We present the basics of the method here; details are in online Appendix C3.

Our overall goal is to simulate the effects of subsidy design in a simple model of insurer price competition (like the model in Section I) using demand and cost parameters estimated from Massachusetts. As in Section I, we model insurers as maximizing static profit, holding fixed the set of competitors and their plans. While these are strong assumptions—e.g., they rule out intertemporal pricing strategies and impacts on product design and entry/exit—they are consistent with other empirical work on insurance markets (e.g., Curto et al. 2014; Decarolis, Polyakova, and Ryan forthcoming; Tebaldi 2017) and let us quantify a benchmark estimate of the relevance of price-linked subsidies.

We specify the following market policies where possible based on ACA rules. We limit the demand sample to consumers above poverty based on ACA subsidy criteria. We also use the mandate penalty and affordable amounts from the ACA, though results are similar if we use CommCare's rules. We use the set of competing plans in CommCare and link subsidies to the cheapest plan since we do not have demand estimates for gold/silver/bronze tiers as in the ACA.[16] We do our analysis

[16] In practice, cross-tier substitution is less important in the ACA because the vast majority of subsidized consumers choose a silver plan in order to get "cost-sharing reduction" (CSR) subsidies linked to silver plans. Another potentially relevant factor is that ACA subsidies are linked to the *second-cheapest* silver plan price (versus the cheapest plan in Massachusetts), which might have lower cross-price elasticity with the outside option.

separately for the CommCare plans in 2009 and 2011 to illustrate the different competitive dynamics before and after CeltiCare's entry in 2010.

Insurers maximize the following profit function:

$$\pi_{jt}(P) = \sum_i \left(\phi_{it} P_{jt} - E(c_{ijt})\right) \cdot Q_{ijt}\left(P^{Cons}(P_t)\right), \tag{8}$$

where $Q_{ijt}(\cdot)$ and $E(c_{ijt})$ are the estimated demand and cost functions, $P^{Cons}(\cdot)$ is the subsidy function mapping prices into consumer premiums, and ϕ_{it} is the consumer risk adjustment score. We model ϕ_{it} as perfectly capturing consumers' expected costs—with the average plan—in a proportional way, which essentially eliminates intensive margin risk selection (see online Appendix C3 for details). Equilibrium is defined by the first-order conditions (FOCs) $d\pi_{jt}/dP_{jt} = 0$ for all j, given other plans' prices and any impact of P_{jt} on the subsidy. With price-linked subsidies, there are sometimes a range of possible equilibria for the cheapest price (see online Appendix C3 for a discussion). When this occurs (for our 2009 simulations), we report results for the minimum and maximum prices consistent with this range.

D. *Simulation Results*

Prices.—Table 4 reports simulation results for 2009 and 2011 market structures, comparing outcomes under price-linked versus fixed subsidies (columns). The first row reports the main outcome of interest: the cheapest plan's price. Price-linked subsidies increase this pivotal plan's price by between \$4 and \$26 (1–6 percent) for 2009 and by \$24 (6 percent) in 2011. The range of estimates for 2009 reflects the fact that in under price-linked equilibrium, Network Health and BMC tie for the cheapest plan with a range of admissible equilibria. In 2011, CeltiCare is the sole cheapest plan, reflecting its much lower costs.

The rest of the table reports additional outcomes, including profit margins, average prices and subsidies, and the share insured. 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 percent to about 41 percent. The ~20 percent 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.

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 percent lower in 2011 and 0.7–7.8 percent 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

Table 4—Equilibrium under Price-Linked and Fixed Subsidies

	2009					2011		
	Price-linked		Fixed	Difference				
	Min	Max		Min	Max	Price-linked	Fixed	Diff
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%
Average price	$418.6	$441.2	$416.6	$2.0	$24.6	$427.3	$405.4	$21.9
Average 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%

Notes: 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 (CeltiCare 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.

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.[17]

The estimated $24 price distortion for 2011 is smaller than the $36 we calculated in Section IIIC using the first-order approximation. We can try to separate out how much of the difference is due to (i) different elasticities in 2008 (when the mandate penalty introduction occurred) versus for CeltiCare in 2011 (the simulation year), (ii) adverse selection and risk adjustment, (iii) pass-through of cost shocks into prices not equaling one (which is implicitly assumed by the first-order approximation), and (iv) strategic responses from the other firms. We analyze the role of each of these factors in online Appendix C4 and online Appendix Table C6. We find that about half of the gap between estimates is explained by different elasticities in 2008 versus 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 one.

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 percent 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 percent. (The lone exception is BMC in the max-price equilibrium in 2009, whose MLR is 77 percent; see table in online Appendix C4).

Welfare Results.—Table 5 shows the different components of welfare—consumer surplus, the avoided externality of uninsurance, government costs, and profits—under fixed and price-linked subsidies. Section IB discusses these components, and online Appendix C3 reports the specifics for calculating their values. The

[17] These 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 percent of the total state population.

Table 5—Costs and Surplus under Price-Linked and Fixed Subsidies

	$ per eligible consumer per month							
	2009					2011		
	Price-linked		Fixed	Difference				
	Min	Max		Min	Max	Price-linked	Fixed	Diff
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

Notes: This is a comparison of welfare under price-linked versus fixed subsidies; see note to Table 4 for additional details. "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.

externality includes both the estimated cost of uncompensated care for the uninsured and a paternalistic social disutility of uninsurance calibrated so that the affordable amounts for price-linked subsidies are optimal. The numbers reported are per *eligible* consumer (including the uninsured) and 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 ~20 percent of the eligible population who purchase CeltiCare.[18]

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 additional person equals the additional public costs.

The last two rows of Table 5 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

[18] We note that the overall level of consumer surplus is relatively small (and sometimes negative) because we compute it relative to a world without the program—i.e., no subsidized insurance *and* no mandate penalty. Because a sizable share of eligible people are uninsured (and therefore hurt by the mandate penalty), it is not surprising that the program has little (or negative) effect on consumer surplus.

may decrease. Even when profits are higher, the sum of public surplus and profits is still lower under price-linked subsidies.

Counterfactual: 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, CeltiCare'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 4. 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 6 shows 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 price distortions range from \$28 to \$50 (or 6–12 percent)—all of which are larger than the \$23.80 we estimate for the full market with five insurers. More notable is how much larger the distortion is when there is a large cost difference between the two plans. For instance, the distortion is only slightly larger (\$28.0) when CeltiCare and Network Health, two low-cost plans, are competing. But it is more than twice as big (\$50.1) when CeltiCare competes against high-cost NHP. Thus, not only the number but the *type* of insurers matters critically for the distortionary effect of price linking: it is less bad when the competing plans have more similar cost structures. This is 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.

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.

V. Discussion

A. *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.

Table 6—Price Distortion with Two Insurers (2011 Parameters)

		(Lower cost) Network Health		⟶ BMC		(Highest cost) 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%)

Notes: (All five insurers in the market: \$23.83 (6.3%).) 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 4. 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.

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 straightforward 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 percent lower and 20 percent 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 online Appendix D we discuss and analyze other reasonable benchmarks, give more details on the uncertainty framework and simulations, 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 percent to $+20$ percent.

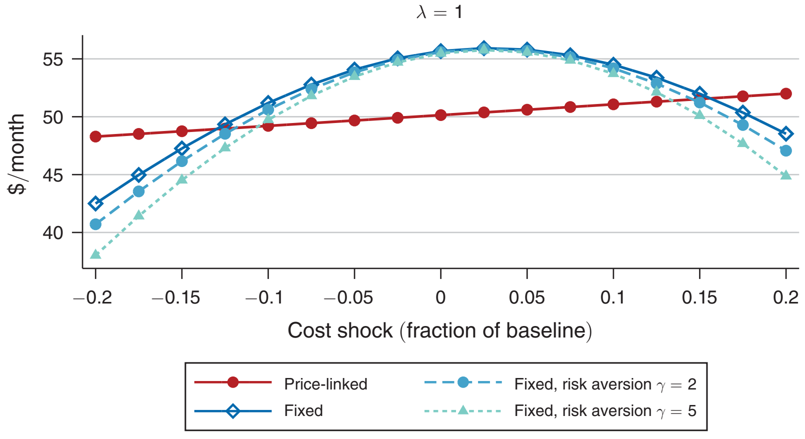

Figure 3. Public Surplus under Cost Shocks in 2011

Notes: 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 percent to +20 percent of baseline. The dashed and dotted lines subtract a cost of pricing risk to consumers (using the method in equation (A2)) with coefficients of relative risk aversion of two and five, respectively.

As expected, based on the lower prices, fixed subsidies result in higher surplus at a cost shock of zero. For nonzero 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 percent or less than −12.5 percent. Using equation (A2), 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 perccent or below −10 percent 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. The range of shocks under which fixed subsidies do better is somewhat smaller, −7.5 percent to +10 percent, but still substantial. These results are also shown in online Appendix D, where we also present results for 2009.

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 US 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 percent, and for 3-year growth it is 4.8 percent. 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 percent (in absolute value) are exceedingly rare. 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 percent—the minimum shock for which price-linked subsidies do better—still occur less than 1 percent 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 are 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, then 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 percent as an index, after

10 years, costs in the NHE data would on *average* have been 9 percent 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.

B. *Implications for Other Markets*

The tradeoffs with price-linked subsidies apply more broadly than the ACA. They apply in any market where (i) firms have market power—i.e., a small price increase does not cause a firm’s demand to fall to zero—and (ii) 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.[19] 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 online Appendix A3) 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’ 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 percent in 2011) (see Decarolis 2015).[20]

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 upward the private school prices. Most housing markets

[19] Medicare 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.

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

have more suppliers, so setting housing subsidies based on market-level housing prices is likely to be less distortionary.[21] Again, price-linked subsidies would only be beneficial if the optimal subsidy increased with prices.

VI. 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 percent 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.

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

Baicker, Katherine, and Douglas Staiger. 2005. "Fiscal Shenanigans, Targeted Federal Health Care Funds, and Patient Mortality." *Quarterly Journal of Economics* 120 (1): 345–86.

Berry, Steven, James Levinsohn, and Ariel Pakes. 2004. "Differentiated Products Demand Systems from a Combination of Micro and Macro Data: The New Car Market." *Journal of Political Economy* 112 (1): 68–105.

Chan, David, and Jonathan Gruber. 2010. "How Sensitive Are Low Income Families to Health Plan Prices?" *American Economic Review* 100 (2): 292–96.

Chandra, Amitabh, Jonathan Gruber, and Robin McKnight. 2011. "The Importance of the Individual Mandate—Evidence from Massachusetts." *New England Journal of Medicine* 364: 293–95.

Chetty, Raj. 2009. "Sufficient Statistics for Welfare Analysis: A Bridge between Structural and Reduced-Form Methods." *Annual Review of Economics* 1: 451–88.

[21] Linking 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.

Chetty, Raj, and Adam Szeidl. 2007. "Consumption Commitments and Risk Preferences." *Quarterly Journal of Economics* 122 (2): 831–77.

Curto, Vilsa, Liran Einav, Jonathan Levin, and Jay Bhattacharya. 2014. "Can Health Insurance Competition Work? Evidence from Medicare Advantage." NBER Working Paper 20818.

Cutler, David M., and Sarah J. Reber. 1998. "Paying for Health Insurance: The Trade-off between Competition and Adverse Selection." *Quarterly Journal of Economics* 113 (2): 433–66.

Decarolis, Francesco. 2015. "Medicare Part D: Are Insurers Gaming the Low Income Subsidy Design?" American Economic Review 105 (4): 1547–80.

Decarolis, Francesco, Maria Polyakova, and Stephen P. Ryan. Forthcoming. "Subsidy Design in Privately-Provided Social Insurance: Lessons from Medicare Part D." *Journal of Political Economy*.

Enthoven, Alain C. 1988. *Theory and Practice of Managed Competition in Health Care Finance*. Amsterdam: North-Holland.

Ericson, Keith Marzilli. 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, Keith Marzilli, and Amanda Starc. 2015. "Pricing Regulation and Imperfect Competition on the Massachusetts Health Insurance Exchange." *Review of Economics and Statistics* 97 (3): 667–82.

Ericson, Keith Marzilli, and Amanda Starc. 2016. "How Product Standardization Affects Choice: Evidence from the Massachusetts Health Insurance Exchange." *Journal of Health Economics* 50: 71–85.

Finkelstein, Amy, Nathaniel Hendren, and Erzo F.P. Luttmer. 2019. "The Value of Medicaid: Interpreting Results from the Oregon Health Insurance Experiment." *Journal of Political Economy* 127(6): 2836–74.

Finkelstein, Amy, Nathaniel Hendren, and Mark Shepard. 2019. "Subsidizing Health Insurance for Low-Income Adults: Evidence from Massachusetts." *American Economic Review* 109 (4): 1530–67.

Garthwaite, Craig, Tal Gross, and Matthew J. Notowidigdo. 2018. "Hospitals as Insurers of Last Resort." *American Economic Journal: Applied Economics* 10 (1): 1–39.

Hackmann, Martin B., Jonathan T. Kolstad, and Amanda E. Kowalski. 2015. "Adverse Selection and an Individual Mandate: When Theory Meets Practice." *American Economic Review* 105 (3): 1030–66.

Handel, Benjamin R. 2013. "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts." *American Economic Review* 103 (7): 2643–82.

Handel, Ben, Igal Hendel, and Michael D Whinston. 2015. "Equilibria in Health Exchanges: Adverse Selection versus Reclassification Risk." *Econometrica* 83 (4): 1261–1313.

Ho, Katherine. 2006. "The Welfare Effects of Restricted Hospital Choice in the US Medical Care Market." *Journal of Applied Econometrics* 21 (7): 1039–79.

Kaiser Family Foundation (KFF). 2018. "Insurer Participation in the 2018 Individual Marketplace." Power Point. Posted November 17. https://www.kff.org/slideshow/insurer-participation-in-the-2018-individual-marketplace/.

Liu, Yiyan, and Ginger Zhe Jin. 2015. "Employer Contribution and Premium Growth in Health Insurance." *Journal of Health Economics* 39: 228–47.

Mahoney, Neale. 2015. "Bankruptcy as Implicit Health Insurance." *American Economic Review* 105 (2): 710–46.

Ruggles, Steven, Katie Genadek, Ronald Goeken, Josiah Grover, and Matthew Sobek. 2015. "Integrated Public Use Microdata Series: Version 6.0 [dataset]." IPUMS, University of Minnesota, Minneapolis, MN. https://ipums.org/projects/ipums-usa/d010.v6.0 (accessed April 25, 2016).

Shepard, Mark. 2016. "Hospital Network Competition and Adverse Selection: Evidence from the Massachusetts Health Insurance Exchange." NBER Working Paper 22600.

Shleifer, Andrei. 1985. "A Theory of Yardstick Competition." *RAND Journal of Economics* 16 (3): 319–27.

Tebaldi, Pietro. 2017. "Estimating Equilibrium in Health Insurance Exchanges: Price Competition and Subsidy Design under the ACA." http://www.pietrotebaldi.com/uploads/8/1/3/6/81364040/eehix_aug14_17.pdf.

Werden, Gregory J. 1996. "A Robust Test for Consumer Welfare Enhancing Mergers among Sellers of Differentiated Products." *Journal of Industrial Economics* 44 (4): 409–13.