NBER WORKING PAPER SERIES

PRICING REGULATION AND IMPERFECT COMPETITION ON THE MASSACHUSETTS HEALTH INSURANCE EXCHANGE

Keith M. Marzilli Ericson
Amanda Starc

Working Paper 18089
http://www.nber.org/papers/w18089

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2012

This paper previously circulated as "Age-Based Heterogeneity and Pricing Regulation on the Massachusetts Health Insurance Exchange." We thank Raj Chetty, David Cutler, Mark Duggan, Jonathan Gruber, Larry Katz, Jon Kolstad, David Laibson, Ariel Pakes, Mark Pauly, Jim Rebitzer, Bob Town, and seminar participants at the University of Pennsylvania, Boston University, and the Southern Economic Association for their thoughtful comments. We acknowledge funding from the Lab for Economic Applications and Policy (LEAP) at Harvard University and the National Science Foundation. Funding was provided by the Lab for Economic Applications and Policy (LEAP) at Harvard University and the National Science Foundation. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Pricing Regulation and Imperfect Competition on the Massachusetts Health Insurance Exchange
Keith M. Marzilli Ericson and Amanda Starc
NBER Working Paper No. 18089
May 2012, Revised October 2012
JEL No. I11,I13

ABSTRACT

We analyze consumer demand and model the effect of pricing regulation under imperfect competition using data from the Massachusetts health insurance exchange. We identify consumer demand using coarse insurer pricing strategies. There is substantial heterogeneity in preferences by consumer type, with younger consumers twice as price sensitive as older consumers. As a result, older consumers face higher markups over costs. Modified community rating links prices for consumers that differ in both costs and preferences. Constrained prices are not simply the population-weighted average of unconstrained prices, because community rating changes the marginal consumer firms face. Tightening rating regulations transfers resources from low cost to high cost consumers, but also reduces firm profits and increases overall consumer surplus. We use our model to examine other insurance regulations. For instance, minimum loss ratios (designed to limit firm profits) will also alter the transfers between consumers. Moreover, risk adjustment will be insufficient to equalize prices across consumer types, as markups still differ. As a result, without a mandate, the market can unravel due to differences in preferences alone

Keith M. Marzilli Ericson
Boston University School of Management
595 Commonwealth Avenue
Boston, MA 02115
and NBER
kericson@bu.edu

Amanda Starc
University of Pennsylvania
astarc@wharton.upenn.edu

An online appendix is available at:
http://www.nber.org/data-appendix/w18089

1 Introduction

Health insurance exchanges (HIEs)–government-run marketplaces for private insurance–raise new questions about the regulation of insurance markets. On exchanges, consumers choose from a set of plans that vary not only in financial characteristics, but also in brand, reputation, and provider network. As a result, insurance plans are imperfect substitutes, giving insurers market power that allows them to charge markups over costs. Insurance markets are also heavily regulated, and these regulations are controversial–for instance, the role of modified community rating and the effect of mandates to purchase a minimum level of insurance coverage. However, there is a lack of research that examines how insurance pricing regulation functions in imperfectly competitive markets.

This paper begins by estimating consumer demand on the Massachusetts HIE, the first HIE established in the United States.[1] Our estimated price elasticities of demand show how insurers' markups over cost can vary by type of consumer. Based on these results, we model the impact of modified community rating regulation, which restricts how prices can vary across consumers within a market. Accounting for imperfect competition is crucial. The traditional analysis of community rating focuses on how the regulation links prices for consumers with different costs, and under the assumption of perfect competition, average prices will match average costs. However, community rating also links prices for consumers with different preferences–particularly, how responsive consumers are to price. These preferences matter for price-setting under imperfect competition. Firms set prices to the marginal, not the average, consumer. Modified community rating changes the marginal consumer and thus the level of markups in the market. Changes in this regulation alter both consumer surplus and firm profits in imperfectly competitive markets.

We use our empirical estimates of consumer demand and our model of pricing under modified community rating to simulate the results of making the community rating regulations more or less strict, and show that stricter regulation raises total consumer surplus, lowers firms' profits, but entails substantial transfers away from lower cost, more price sensitive, younger consumers. We then show how modified community pricing regulation interacts with other important and controversial insurance market regulations: minimum loss ratios (which attempt to limit insurer profits), risk adjustment (which attempts to equalize insurers' costs from different groups), and mandated insurance purchase (which attempts to ensure market participation). Our results show that risk adjustment based on costs alone is not sufficient to equal prices across demographic groups, because price sensitivity and

[1] See also Frank and Lamiraud (2009) on the health insurance market in Switzerland, which bears many similarities to the ACA's exchanges.

markups vary across groups. Our simulations show that if the mandate is removed, markets can unravel–even with perfect risk adjustment– due to differences in preferences alone if participation elasticities are large enough.

HIEs are an ideal context to study these issues, as they offer a wide range of choice to consumers. HIE are also interesting and important in themselves. The 2010 Affordable Care Act (ACA) calls upon states and the federal government to set up HIEs to facilitate the purchase of insurance, and a projected 20 million individuals will purchase through the exchanges (CBO 2012). HIE are novel regulatory environments, and states will have substantial latitude in designing and regulating these exchanges. Their choices will shape the market for individually-purchased health insurance. For instance, HIEs can control what types of insurance plans are offered, how information about them is presented to consumers, how prices can vary across consumers, the defaults individuals face, and the frequency of the open enrollment period. However, little is known about the nature of demand for health insurance in such a setting. Understanding consumer demand and insurer incentives is important for both exchange design and for the broader regulation of insurance markets.

The Massachusetts HIE provides an early look at a comprehensive HIE in action, as it is similar to the ACA's exchanges, but was created by Massachusetts' own reforms and has been providing coverage since 2007. While regulation of HIEs is still in flux, HIEs differ in important ways from the way most individuals in the U.S. currently obtain health insurance: employer-sponsored insurance or government-provided coverage (i.e. Medicare or Medicaid). Other markets also offer insight into HIEs, but have crucial differences. Medicare Part D's insurance exchange has been a fruitful field of research,[2] but offers a limited type of coverage (prescription drug insurance) to a narrow age range (the elderly). Employer-sponsored insurance typically offers a limited range of choice (see Dafny, Ho, and Varela [2010] on the value of choice in these contexts), though some large employers may offer a range of choice akin to an HIE. However, employer-sponsored insurance differs from HIEs in how it is regulated (e.g. plans cannot price differentially by age in employer plans) and in the nature of competition (the employer negotiates directly with insurers). Finally, individual markets for direct purchase of insurance are currently small and decentralized (covering only about 5% of the population), with high search costs (Maestas, Schroeder, and Goldman [2009]) and limited public data.

Our analysis of consumer choice on the Massachusetts HIE provides the basis for counterfactual simulations that examine how modified community rating regulation interacts

[2]Abaluck and Gruber (2011) find systematic mistakes in choice, Ericson (2012) shows the role of inertia in firm pricing, and Ketcham et al. (forthcoming) argue that consumers learn over time. See Duggan, Healy and Scott Morton (2008) for more detail on the Medicare Part D reform.

with insurer incentives under imperfect competition. We build on the existing literature that shows evidence of insurer market power in other contexts (see Dafny, Duggan, Ramanarayanan [forthcoming], Dafny [2010], and Starc [2010] for examples). We knit together this literature with the literature on community rating laws, which has examined the effect of community rating on coverage rates (Simon [2005], Zuckerman and Rajan [1999]), and equilibrium outcomes assuming perfect competition (see Finkelstein, Poterba, and Rothschild [2009] on gender-based pricing in annuities, Geruso [2011] on preference heterogeneity under perfect competition). Our work, however, emphasizes the importance of accounting for imperfect competition when analyzing these regulations.

We focus on a particular form of modified community rating: age-based pricing regulation, which plays a major role in HIEs, but has been relatively neglected by researchers. Under this regulation, older individuals pay higher prices than younger individuals only within a certain band, up to the maximum allowable price ratio (MAPR). In Massachusetts, this ratio is 2, so the oldest enrollees cannot be more than twice the price charged to the youngest. Age-based pricing regulation has a substantial effect on the prices individuals face. Limits on age-based pricing are a type of modified community rating, which restricts how insurers can vary prices based on consumer characteristics. Pure community rating, in which all consumers in a risk pool face the same price, creates a trade-off: welfare losses from adverse selection, against which are weighed welfare gains from insuring consumers against the possibility of being a bad risk by having higher expected medical spending. Modified community rating, in which insurers are allowed to vary premiums across consumers–but only within limits–attempts to mitigate some of the welfare loss.

We estimate a discrete choice model of consumer demand on the Massachusetts HIE, focusing on how price sensitivity varies by age.[3] We use discontinuities in the ways firms set their prices by age to overcome the endogeneity problem and identify consumers' response to price.[4] There is substantial variation by age in demand elasticities: younger consumers are more than twice as price sensitive than their older counterparts. Price-sensitivity varies by age, thus insurers would want to price discriminate and charge higher prices to older enrollees, even if health costs did not vary by age. In contrast, existing work has assumed that insurers price differentially by age solely due to cost differentials (e.g., Blumberg, Buettgens, and Garrett 2009), so assumes that regulations only bind to the extent the ratio of costs

[3] Existing estimates of price sensitivity for health insurance vary substantially by context, and most examine employer-sponsored health insurance. The only work addressing HIEs specifically is Ericson and Starc (2012), whose measures of price sensitivity do not account for the endogeneity of premiums; instead, that paper focuses on the role of heuristics in choice. See also Carlin and Town (2007), Gruber and Washington (2006), Bundorf, Levin and Mahoney (2008) and Einav, Finkelstein and Cullen (2010)..

[4] Plan prices are constant within five year age blocks (e.g. 30 to 34), and then jump.

exceeds the maximum allowable ratio of prices by age. In general, though, profit maximizing insurers will price discriminate by responding to heterogeneity in consumer demand that is correlated with observable tags. Here we study age, but similar rationales apply to family size, gender, and geographic location.

Using our estimates, we simulate insurers' optimal pricing under various models of market structure and examine the potential welfare effects of alternative pricing regulations. We conduct a counterfactual exercise in which estimate the distributional consequences of eliminating or tightening age-based pricing rules. Accounting for price discrimination is crucial in estimating the effect of such regulation. Compared to unconstrained insurer prices, the Massachusetts maximum allowable price ratio of 2 is estimated to raise prices on young consumers by 22%. However, even constraining the ratio of prices by age to ratio of costs by age (i.e. a maximum allowable price ratio of 2.75, near that of the ACA's ratio of 3) still leads to price increases on younger consumers of 7%; these increases would not occur under perfect competition.

Despite these transfers[5], modified community rating improves surplus overall. The imperfectly competitive insurers price to the marginal, rather than the average, consumer. The marginal consumer in a plan is likely to be a young consumer with a relatively high price elasticity and relatively low medical expenditures. As a result, with binding modified community rating, average prices, profits, and markups are lower than without regulation. Moreover, minimum loss ratios, which cap insurer profits, maintain competitive pressure on insurers while dampening transfers between groups, since they reduce insurers' incentives to set high prices for price-insensitive older consumers.

Finally, modified community rating can lead to severe welfare losses in the absences of an effective mandate. Even if risk-adjustment were perfect, so that insurer costs did not differ by enrollee age, differences in preferences alone can lead to this market unraveling. If consumers are allowed to opt out of coverage, modified community rating restrictions will lead the most price-sensitive consumers opt out. As these consumers opt out, less price-sensitive consumers are left in the market, leading to higher markups. This, in turn, leads more price-sensitive consumers to opt out of the market. If the participation elasticities are high enough, a death spiral can result (Cutler and Reber 1998) where there are only price insensitive consumers left in the market. As a result, a weak or absent mandate may negate the consumer surplus gains achieved from modified community rating.

The preference heterogeneity that we identify is critical for analyzing policy, and this implies that choosing the set of consumers who form a risk pool is critical for determin-

[5] We note that it is not clear why one would want to transfer from younger to older consumers in this way.

ing the functioning of and allocation of surplus within insurance markets. When defining which segments of consumers to include in a health insurance exchangem such subsidized enrollees, younger consumers eligible for catastrophic plans, or employees in small groups,. understanding differences in preferences is as critical as understanding differences in cost.

The paper proceeds as follows. Section 2 describes the Massachusetts HIE, rating regulation, and some reduced-form results. Section 3 details our identification strategy, and develops and estimates our model of consumer demand. Section 4 derives optimal pricing behavior under imperfect competition and modified community rating regulation. Section 5 simulates insurance market outcomes under alternative age-based pricing regulations and examines the effect of minimum loss ratios. Section 6 describes the effect of removing insurance mandates. Section 7 concludes.

2 The Massachusetts Connector: Context and Data

2.1 Massachusetts' Health Reform

The state of Massachusetts signed its health care bill into law in April 2006, with the goal of providing universal coverage for its residents; the reform, in turn, served as a national model for health reform. This reform had many features, including expansions in public coverage, and individual and employer mandates. A key feature of this reform was the individual mandate, which required all Massachusetts residents to purchase a minimal level of health insurance coverage (minimum creditable coverage),. or face a penalty equal to half of the premium of the lowest cost health insurance plan offered through the exchange. To facilitate consumers purchasing insurance, the state required employers with 11 or more employees to make a fair and reasonable contribution to employees' health insurance costs. It also established the Commonwealth Care Program, which provided free or subsidized coverage to lower income residents, who earned up to 300% of the federal poverty level.[6]

Finally, the reform established an unsubsidized health insurance exchange that is the focus of this paper. The Massachusetts HIE (known as the Commonwealth Choice Program) was designed to facilitate nongroup coverage purchased directly by households and small group purchase of insurance. It has offered health insurance since May 1, 2007 (with the mandate taking effect July 1, 2007). Both the HIE and the Commonwealth Care Program are run by the Commonwealth Connector Authority (the "Connector"), a quasi-public agency that shapes the market for individual coverage in Massachusetts in a number of ways. The

[6] The Commonwealth Care program is quite different from the HIE established by the ACA, as enrollees do not choose among plans of different generosities, but are assigned to a plan tier based on their income. Additionally, prices in Commonwealth Care do not vary by age.

Connector operates the exchange's website,[7] actively negotiates with insurers, and chooses which features of insurance plans are highlighted.

After the Massachusetts reform, the rate of uninsurance dropped. By 2009, 97.3% of the population was insured (Long and Phadera 2009); the insurance rate for residents above 300% of the poverty line (i.e. those eligible for the exchange) was 99.1%. Increases in the insured came from individuals purchasing insurance through the Massachusetts HIE, through increased offering of employer-provided health insurance, and through expansions in subsidized coverage (Gruber 2011). Kolstad and Kowalski (2010) show that the Massachusetts reform not only increased coverage, but also decreased hospitalization for preventable conditions. However, the effect of the health reform and the HIE on the level and growth rate of premiums is a point of contention. By characterizing consumer demand and market structure in this paper, we provide a foundation for future analysis of the impact of the reform.[8]

2.2 Regulation of the Health Insurance Exchange

In addition to the individual mandate to purchase insurance, there are three other important regulations for the HIE that we consider:

Minimum creditable coverage (MCC): MCC is the least generous plan that is sufficient to comply with the mandate. The Connector is responsible for determining MCC for the state based on a combination of actuarial value, out-of-pocket maximum, deductibles, covered physician visits, and prescription coverage. In Massachusetts, MCC includes prescription drug coverage and three check-ups, caps deductibles at $2000 for an individual and $4000 for a family, and caps out-of-pocket expenditures at $5000 for an individual and $10,000 for a family. A large number of policies just satisfying MCC are available, and they are quite popular. Therefore, regulation regarding the definition of MCC is likely to be important in a market with a mandate.[9]

Modified Community Rating: Modified community rating rules apply to pricing on the exchange. Prices for products can vary by age (and geography), but the ratio of a price for a given product for any two individuals cannot exceed the maximum allowable price ratio, which is 2 in Massachusetts. This establishes an age-band within which prices can vary. In addition, no medical underwriting is allowed, and plans are guaranteed issue (no one can be denied coverage). These rating rules are critical in shaping premiums in the market. Age, in particular,. is a critical feature of pricing.

[7] The website is http://www.mahealthconnector.org.

[8] For an overview of the Massachusetts reform and its various components, see Ericson and Starc (forthcoming), Gruber and Levy (forthcoming), Miller (forthcoming), and Long (forthcoming).

[9] Finkelstein (2004) finds that minimum standards can reduce enrollment by potentially exacerbating adverse selection. However, in the presence of a mandate, such concerns are much less pressing.

Figure 1 shows that age-based pricing regulations are binding: the average monthly premium for a 27-year-olds is just over \$300, which the premiums for older consumers are just over \$600. While the ACA specifies a 3:1 maximum allowable age rating band in the individual health insurance exchanges, states can impose more strict regulation. For example, Maryland has chosen a price ratio of 2.8 (Carey and Gruber 2010). The choice of this ratio will alter the both the size and division of surplus among young consumers, older consumers, and firms. We note that price discrimination through plan design (second degree) is largely prohibited by the mechanisms in place by the regulator, especially minimum creditable coverage, which outlaws catastrophe plans with extremely high deductibles (i.e. \$10,000/year).

Minimum Loss Ratios: Minimum (medical) loss ratios require that insurers pay out a certain percentage of premiums in medical claims. The Massachusetts reforms did *not* contain minimum loss ratio regulation, and no such regulations were in effect during the time period we analyze. However, the ACA requires a minimum loss ratio of 0.80 for individual/small group markets and 0.85 for large employers beginning in 2011.

2.3 Making Choices on the Exchange

The exchange offers a variety of health plans administered by the major private insurers in the state.[10] Insurers had relatively wide latitude in designing these plans, which were grouped into tiers based on actuarial value: bronze, silver, and gold.[11] Bronze plans are generally less generous (higher cost-sharing) and therefore tend to be cheaper. Gold plans are the most generous, and hence most expensive, while silver plans forge a middle ground. In addition to this main market, there is a separate market for young adult consumers aged 18 to 26, in which plans tend to have more limited coverage, such as optional prescription drug coverage.

Consumers face a number of steps when purchasing insurance from the exchange. (Screenshots from the purchasing process are included in the Appendix.) After entering demographic information, consumers are offered a choice of plans that vary along a number of dimensions, including copayments, deductibles, and premiums. Importantly, the plans are placed into tiers; this grouping might affect consumer choice. Even within each tier, consumers must

[10] In our sample, the following firms sold insurance via the connector: Blue Cross Blue Shield of Massachusetts, Fallon Community Health Plan, Harvard Pilgrim Health Care, Health New England, Neighborhood Health Plan, and Tufts Health Plan. All of these insurers are non-profit, but Dafny and Ramanarayanan (2012) show that for-profit and not-for-profit insurers behave similarly.

[11] Note the ACA requires that HIEs contain four tiers (it adds platinum), with slightly different definitions for the tiers. Also, beginning in 2010 (after our sample), the Connector required plans to take one of six standardized forms (bronze low, bronze high, etc.), though plans may still differentiate themselves based on their provider networks.

weigh multiple dimensions of plans (copayments, coinsurance, dental coverage).[12] Finally, the consumers enroll.

The website itself, in addition to regulation, has the potential to shape consumer choices. Tiering can also affect how insurers design plans; for example, they may design plans to meet the minimum level of generosity in a tier. The way information is presented, plan features highlighted, and the order in which plans are sorted, may also affect consumer behavior. For instance, Ericson and Starc (2012) finds a discontinuity in preference for the minimum choice plan. During the initial period, the website sorted plans according to price (as opposed to, for example, consumer satisfaction) so consumers may have inferred that price was the most important variable differentiating these plans.

2.4 Data and Descriptive Statistics

We use transaction-level data (purchase, cancellation, and payments) from the unsubsidized market (Commonwealth Choice) from its launch in July 2007, until December 2009. We observe approximately 50,000 transactions. There are large spikes in initial enrollment during the first month of the exchange's operation, as well as just before the individual mandate's financial penalties took effect in December 2007, with a steady-state enrollment of approximately 1,000 households per month. Appendix Figure A.1 plots a histogram of the number of individuals choosing single coverage joining the Connector for the first time, by month (the majority of purchases are for single coverage).

Table 1 describes the demographics of these consumers: most are young, with an average age of approximately 35. Most purchase individual, rather than family, plans, and a sizable percentage lives in Middlesex County (which includes Boston suburbs like Cambridge and Somerville). The average premium paid is about $420.30 per month, but varies substantially by age and plan tier. The median consumer tenure in the HIE is 1.25 years, with only small differences between tier of plan chosen. For more detail on duration of enrollment in the exchange, see Ericson and Starc (forthcoming).

We limit our sample to enrollees purchasing coverage for the first time from the exchange,[13] and focus on consumers purchasing individual coverage (as opposed to household coverage), since the majority of plans sold are of this type. We exclude consumers eligible for young adult insurance (those aged 26 and under) from this sample because they have a different choice set and rarely purchase plans other than young adult insurance plans. Our choice analyses focus on a subset of the data: November-December 2009. Because we observe

[12] Beginning in 2010 (after our sample period), the financial characteristics of plans were standardized in each subtier.

[13] We exclude repeat purchases since they may display inertia (Ericson 2012; Handel 2009).

transaction-level data, we do not observe all the plan prices that individuals face. However, for November and December 2009, we collected an extensive set of price quotes from the exchange's website using a Perl script, allowing us to create the plan menu faced by each enrollee. The Data Appendix gives more details. We are also able to run robustness checks on a subset of the data that extends back until July 2009 (see Data Appendix). The choice of sample period does not have an effect on our results.

3 Model of Consumer Demand

3.1 Theoretical Model of Demand

We model consumers' choice of insurance plan using a standard discrete-choice logit model. We assume that consumer i's utility of plan j in market m is given by:

$$u_{ijm} = \delta_{jm} + \mu_{ijm} + \varepsilon_{ijm},$$

where δ_{jm} is the mean utility of a plan in market m, μ_{ijm} represents the (mean-zero) component of a plan's utility that varies based on observed individual characteristics (e.g., age), and ε_{ijm} is an error term that is independently and identically distributed (i.i.d.) extreme value. This implies shares can be written as:

$$s_{ijm} = \frac{\exp\left(\delta_{jm} + \mu_{ijm}\right)}{1 + \Sigma_j \exp\left(\delta_{jm} + \mu_{ijm}\right)},$$

where s_{ijm} represents the probability that consumer i purchases product j in market m. In the absence of individual heterogeneity μ_{ijm}, the δ_{jm} parameters simply represent an inversion of the observed market shares for each plan. The mean utility δ_{jm} can be decomposed at price into price and plan characteristics, where $\delta_{jm} = \alpha p_{jm} + X'_{jm}\beta + \zeta_{jm}$, with X_j being either a vector of plan fixed effects or a vector of plan characteristics, such as insurer (brand), deductibles, and copayments or, alternatively, plan fixed effects, p_{jm} is the price, and ζ_{jm} is any unobserved product characteristic.[14] A given insurance plan (e.g., HMO Blue Basic Value) is offered to all ages 27-65, and in multiple (but not necessarily all) geographical markets.

[14] When plan fixed effects are included, this is the market specific deviation from the mean, as in Nevo (2001).

3.2 Identification Strategy: Coarse Firm Pricing By Age

The level of consumer price sensitivity is an important feature of demand that affects both policy design and insurer price-setting behavior. Estimates of price sensitivity are difficult to identify because unobserved plan characteristics may be correlated with price. We use a discontinuity identification strategy based on coarse pricing rules used by firms to identify the effect of price on choice. By law, firms may vary prices continuously (within broad limits) by age. However, firms do not, in fact, set a different price for each age: they price more coarsely. The listed premium given plan has discrete changes at ages ending in 0 or 5 (30, 35, 40, etc.). Figure 1 shows jumps at each age in the average premium for a constant set of plans.[15] These jumps translate into very similar consumers facing very different vectors of premiums: the underlying preferences of a 39-year-old and 40-year-old will likely be very similar; however, they will face different premiums for the same plan. Appendix Section A.2 runs regressions that show exactly how much prices change at each discontinuity. It also motivates the discrete choice model by showing that reduced form regressions of total spending on price will be misleading if they do not take into account that fact that consumers are constrained (about half of enrollees are already choosing the least generous tier of plan).

Our identification strategy relies on the discontinuity in price being unrelated to demand; i.e. preferences evolve continuously as an individual ages, so turning 40 is like turning 39 or 38. Based on our conversations with insurance firms, these discrete jumps in price result from firms' menu costs when setting premiums. The combination of age bins and zip codes alone gives rise to over 40,000 potential prices, all of which must be submitted to a regulator for approval. Moreover, Chu, Leslie, and Sorenson (2011) show that firms can obtain profits close to a perfectly price discriminating firm using coarse pricing rules. Finally, while insurers price coarsely in this time period, more recent data (August 2012) show that some firms have begun to allow prices to vary continuously by age. This supports our identification strategy; that demand and costs do not jump at round-numbered ages.

While firms could price in discrete age blocks if the cost of an insured individual changed dramatically at each age cutoff, this alternative explanation for the jump in prices is not supported by the data. While diagnostic tests (such as mammograms) are recommended for patients beginning at the age cutoffs, observed medical spending in the Medical Expenditure Panel Survey (MEPS) rises smoothly and shows no systematic discontinuities in health expenditures at round numbered ages. Thus, differences in spending are unlikely to account for such large price jumps.

[15] The marginal cost of choosing a more generous plan jumps correspondingly, as shown by Appendix Figure A.2. The ratio of the cost of the average gold plan to that of the average bronze plan varies slightly within each age category but stays between 1.8 and 2.

We identify the premium coefficient α under the assumption that preferences evolve continuously with age, so that discontinuities in mean utilities at round-numbered ages are solely attributable to discontinuous changes in premiums.[16] More formally, let $\delta_j + \mu_{j30}$ be the mean utility of product j offered to consumers who are age 30, and $\delta_j + \mu_{j29}$ be the mean utility of product j offered to a consumer who is age 29. Our choice model gives us consistent estimates of both utilities. Then in the absence of age trends, the price coefficient can be simply written as:

$$\alpha = E_j \left[\frac{\mu_{j30} - \mu_{j29}}{p_{j30} - p_{j29}} \right].$$

Of course, there are age trends in preferences. We allow for u_{ij} to evolve continuously over ages, but limited data require that we place some structure on how it does so. Estimating a separate linear spline in age for each of the 21 different plans would allow a great deal flexibility but is infeasible in our data. We allow for preferences for plan tier (bronze, silver, gold) to evolve flexibly with age. Further, we allow different plans within a tier to have different qualities. In other specifications, we allow a linear trend in the age-plan fixed effect interaction. Ultimately, the assumption that identifies the price coefficient in these specifications is that age-specific deviations in preference for plans within a given tier are not correlated with prices. This seems a reasonable assumption, and our results do not change substantially when we allow for more variability in age-specific preference for plans.

We use coarse geographical pricing to define markets. Instead of varying prices for each zip code, firms set prices for larger geographic regions that roughly correspond to hospital referral networks; these may be a good proxy for underlying insurer costs. For example, Blue Cross Blue Shield charges three sets of premiums: one set for western Massachusetts, one set for the greater Boston area, and one set for Cape Cod. We use this variation to define a geographical region that is a set of zip codes in which prices do not vary within a plan-age cell. (See Data Appendix for details.)

Appendix Table A.2 supports our identification strategy by showing that characteristics of enrollees' zip codes do not change discontinuously between age categories, with the exception that enrollees over age 55 seem to be slightly more wealthy, employed, and white. This may lead us to slightly underestimate the price sensitivity of this age category. Similarly, the density of individuals enrolling in the exchange does not change at the various age cutoffs. Figure A.3 shows the number of enrollees in each one-year age bin (we do not have exact

[16] Without assuming that preferences evolve smoothly, the premium coefficient is not separately, nonparametrically identified from variation in preferences over plans. This is because the premium is itself a (highly nonlinear) function of demographic characteristics, such as age, that may also impact preference for plans or benefit designs. This is why allowing flexible preferences over plans is crucial for correctly identifying price sensitivity.

birthdate, only age in years). Visual inspection indicates there is no general pattern of densities dropping at round numbered ages, with perhaps an anomalous low enrollment for individuals aged exactly 50 years. The final column of Appendix Table A.2 shows that the density doesn't change discontinuously at any breakpoints, with the potential exception of the age 50 breakpoint.[17] In reduced-form analyses (available in the appendix), we find that total spending rises nearly one-for-one with price increases. However, this misleadingly suggests little consumer response to price. Consumers are already clustered at the cheapest plans (25% choose the cheapest plan available to them in our sample period), they have little ability to reduce insurance spending in response to price increases. A discrete choice model is therefore more appropriate in measuring the value that consumers place on plans, as it takes into account their limited potential range of substitution.

3.3 Estimation of Consumer Demand

The model is estimated using a conditional logit approach in Table 2. This allows us to interact consumer characteristics, such as age, with plan characteristics, such as price. Panel A, Column 1 estimates the model without allowing for consumer heterogeneity, while column 2 allows price sensitivity to vary linearly by age. The data strongly reject constant price sensitivity by age. Furthermore, by comparing columns 2 and 3 we see that accounting for variation in preferences is important to estimating the level of price sensitivity as well. The estimates in column 3 indicate that the oldest consumer in our sample (64) is roughly half as price sensitive as the youngest consumer in our sample.

Alternative specifications confirm the pattern of lower price sensitivity by age. Panel B of Table 2 separately estimates the model by five-year age bands. These results for each 10-year age bin still show the variation in price sensitivity by age, indicating our structural assumptions are not too restrictive. Finally, for use in some counterfactual exercises, Panel C of Table 2 divides the age span in the exchange in half, and runs the model separately for those under age 45 and age 45 and older. The younger consumers are substantially more price sensitive, and the difference is statistically significant. Furthermore, the results indicate semi-elasticities (which describe the percent change in enrollment given a $100 increase in premiums) for younger consumers of around -3, and for older consumers of just above -1. Figure 2 maps out these price elasticities under various specifications, highlighting the pattern in semi-elasticities over the life cycle.

Table 3 shows additional specifications. The first two specifications break out price sensitivity into five-year age bands. Column 2, in particular, shows that the trend may

[17] The results using the expanded July-December 2009 robustness check sample are similar and, in fact, do not contain any significant differences in zipcode characteristics at age 55.

not be linear, but that the oldest consumers have substantially different preferences than their younger counterparts. Column 3 estimates a nested logit, in which consumers first choose a plan tier (gold, silver, or bronze), and then choose a policy from within that tier. The dissimilarity parameter is an inverse measure of the correlation between error terms in each nest, and a dissimilarity parameter of one would indicate the model collapses to the conditional logit. The low dissimilarity parameter for bronze and silver plans indicates that consumers see these plans as fairly close substitutes. However, with an estimated dissimilarity parameter near 1, gold plans are not close substitutes, indicating that networks and brand name factor highly in the decisions of consumers who are likely to purchase gold plans. Column 4 includes mixed logit results, in which the price coefficient α is allowed to take on a log-normal distribution, shifted by age category. The results show that distribution is shifted toward zero (less price sensitive) for older consumers. Nonetheless, because of data limitations and the flexibility the specifications provide, these results are somewhat noisy. Additional robustness checks that replicate the results for a smaller "bandwidth" are available in Appendix Table A.3.

Taken together, the results from Tables 2 and 3 are striking. The elasticities of the youngest group, those 27 to 35, are nearly twice as large in magnitude as those of the oldest group. The raw data driving these results can be summarized as follows: For older consumers, the marginal cost of gold plans relative to bronze plans is much higher than for younger consumers (Appendix Figure A.2). Despite these differences, the fraction of consumers purchasing bronze plans stays relatively flat with age. This indicates that older consumers have a lower distaste for price relative to their preference for more generous coverage.

Various demographic factors could be driving the preference heterogeneity that we see in the data. For the pricing exercise in the next section, it does not necessarily matter whether age is simply a signal for another demographic factor correlated with preferences or not–insurers can price differently by age, but not on other factors correlated with age (such as income). We note that younger consumers are not from lower income zipcodes in our data. However, because older individuals are more likely to be married, the selection of older consumers into the exchange may differ, as some married consumers have access to insurance through a spouse. In addition, older consumers are less likely to report that they are in excellent health; Strombom et al. (2002) report that older and sicker consumers tend to be less price sensitive. Finally, the relatively older consumers in our sample might be more financially sophisticated, leading them to more heavily weigh characteristics other than price when making decisions.

3.4 Accounting for Selection into the Exchange

While our data describes purchase conditional on participation in the exchange, we know very little about the behavior of consumers outside the exchange. We address selection in a number of ways. First, we discuss in detail what alternative options consumers have besides the exchange, and how this may affect our estimates. Second, we run robustness checks in our model of consumer demand to ensure that our results are not sensitive to the exact definition of the outside good. Furthermore, Section 6 explores the impact of non-participation in the HIE through additional simulations.

Potential enrollees (those above 300% of the poverty line) in the Connector have few close substitutes. If they (or their spouse) are offered employer-sponsored insurance, they are not eligible to enroll in the exchange. Their alternative to the exchange plans are either 1) remain uninsured, or 2) purchase an individual plan from a broker or directly from the insurer. Uninsurance among this group is extremely low (compliance with the mandate is extremely high.) The uninsurance rate is only 0.9% for Massachusetts residents above 300% of the poverty line (calculated from Phadera and Long's [2010] detailed tabulations of 2010 Massachusetts Health Insurance Survey), suggesting very limited scope for consumers to substitute into the HIE if insurer prices were lowered.

The second alternative is individual plans not purchased from the exchange, which we do not have data on. However, we can compare whether there is differential selection based on age into the exchange versus the individual market. The American Community Survey (ACS) provides information on the number of individuals who purchase insurance *directly from the insurance company*. The 2008-2010 ACS for Massachusetts shows that age distribution of people who directly purchase is relatively flat with respect to age (See Appendix Table A.4), with slightly more under 30 year olds. This pattern is quite similar to the age distribution of individuals in the exchange (see Figure A.3), indicating that there is no systematic selection into the HIE by age.

In addition to comparing our data to other sources, we can also do a number of robustness checks to ensure that the results are not sensitive to the exact specification of an outside option. Table A.5 shows that the difference in price sensitivity by age is robust to assuming gold and silver plans to be the outside option and identifying off variation in bronze prices alone. As a result, selection into the exchange has little effect on the pattern of our demand estimates. By contrast, in the absence of mandate, accounting for non-participation in the market would be quite important; this is addressed in Section 6.

4 Firm Pricing and Age-Based Pricing Regulation

4.1 Motivation

In this section, we model how age-based pricing regulation affects markets in the presence of imperfect competition and age-based heterogeneity in price sensitivity. We develop the model in the context of age, but the same logic would apply to any observable tag by which costs and preferences varied.

We analyze the effect of three types of age-based pricing regulations:

- Age-Pooling: firms cannot vary prices by age.
- Age-Bands: firms can vary prices by age, but with a maximum allowable price ratio of θ (the ratio of the highest price to lowest price).
- Age-Unconstrained: firms can vary prices by age.

In all cases, we assume that if a plan is offered, it must be offered to all ages. Note further that pooling and unconstrained prices are simply special cases of Age-Bands (where θ =1 and ∞, respectively). On the Massachusetts HIE, $\theta = 2$, while the ACA requires states to set $\theta \leq 3$.

4.2 Model and Theoretical Predictions

Consider two types of consumers, old and young, who are purchasing an insurance plan.[18] Costs rise with age, so the old have an average cost (to the insurer) of c_H, greater than the cost of the young c_L. Let fraction σ of the population be old, and fraction $1 - \sigma$ be young. There are $N \geq 2$ profit-maximizing insurers, each offering a single plan[19] that is available to the young and the old. Insurers can determine whether an individual is old or young, but cannot further determine the expected cost of the individual. Hence, each insurer can set two prices, one for old and one for young individuals: p_H and p_L.

We first examine how regulation affects pricing in perfectly competitive markets, in which products are identical and firms make zero profits.[20] Stricter limitations on age-based pricing transfers resources from old individuals to young individuals in addition to changing the level

[18] Here, we assume both groups have a high enough willingness to pay to purchase insurance, so that selection out of the market is not an issue (i.e., that the mandate is effective). Section 6 examines the consequences of selection with a weak or absent mandate.

[19] We abstract away from adverse selection between policies of different quality.

[20] Formally, let there be a continuum of consumers normalized to measure 1. When multiple firms offer a plan at the same price, consumers are evenly distributed across the firms.

of prices. Age-bands are only binding up to the ratio of costs between the two groups. Prices under perfect competition are summarized below:

- Under the Age-Pooling regulation, prices are equal to population average cost: $\bar{p} = \sigma c_H + (1-\sigma) c_L$.
- Under the Age-Unconstrained regulation, prices are equal to each type's average cost: $p_H = c_H$ and $p_L = c_L$.
- Under the Age-Bands regulation with $\theta \leq \frac{c_H}{c_L}$, prices for the young are above their cost, and for the old are below their cost: $p_H^* = \theta p_L^*$ and $p_L^* = \frac{1}{(1-(1-\theta)\sigma)} \left[\sigma c_H + (1-\sigma) c_L\right]$. When $\theta > \frac{c_H}{c_L}$, the regulation does not bind, and so $p_H^* = c_H$ and $p_L^* = c_L$.

However, in an imperfectly competitive market, the prices set by insurers for each group are determined not only by costs, but also by that group's elasticity of demand. Hence, prices for old and young consumers may differ due to a *price discrimination motive* as well as a *cost differential motive.* Thus, we must consider how characteristics other than cost affect prices when modeling age-based pricing regulations. Price discrimination can amplify price differences if high-cost consumers have lower price sensitivities, while it can reduce (or even reverse the direction of) price differences if low-cost consumers are less price sensitive (the "worried-well"; see Starc [2012]).

Now let the market be imperfectly competitive. Let $\tilde{s}_{ja}$ reflect the share of age group a that purchases insurance at firm j, and let $s_{jH} = \sigma \tilde{s}_{jH}$ and $s_{jL} = (1-\sigma)\tilde{s}_{jL}$ be the number of each group purchasing insurance at firm j. Then, we can write the profits of firm j as:

$$\Pi_j = s_{jH}(p_{jH} - c_H) + s_{jL}(p_{jL} - c_L).$$

Firms set prices based on their first-order conditions (which we assume are unique), subject to the age-based pricing regulations they face. We drop the j subscripts below. We define a few terms: let s_H and s_L be functions of p_H and p_L, respectively, so that s_i' gives the change in type i's enrollment as p_i changes. Let total enrollment be $S = s_H + s_L$. For use in the Age-Bands pricing, define weighted demand $\bar{S} = \theta s_H + s_L$. When the bands are binding, write p_H as an implicit function of p_L, and $\bar{S}$ as a function of p_L, so that $\frac{d\bar{S}}{dp_L} = \theta^2 s_H' + s_L'$.

Proposition 1 *Assume markets are imperfectly competitive. Then, under the Age-Pooling regime,* $p^{Pool} = \frac{1}{\frac{dS}{dp}}\left(s'_H c_H + s'_L c_L\right) - \frac{S}{\frac{dS}{dp}}$. *Under Age-Unconstrained,* $p_H^{Un} = c_H - \frac{s_H}{s'_H}$ *and* $p_L^{Un} = c_L - \frac{s_L}{s'_L}$. *If Age-Bands are binding,* $p_L^{Band} = \left(\frac{\theta s'_H}{\frac{d\bar{S}}{dp_L}} c_H + \frac{s'_L}{\frac{d\bar{S}}{dp_L}} c_L\right) - \frac{\bar{S}}{\frac{d\bar{S}}{dp_L}}$ *and* $p_H^{Band} = \theta p_L^{Band}$.

Proof. Immediate from first-order condition. ■

Proposition 1 shows that under the Age-Unconstrained policy, firms simply set prices for each group equal to cost, plus a markup inversely proportional to the elasticity of that group's demand. An insurer can only set one price under Age-Pooling, which is equal to a markup term inversely related to the elasticity of population demand, plus a cost term, where the relative weight on each cost term is that groups' share of the marginal change in demand. Note that firms price to marginal cost, not average cost: more weight will be put a group that is more price sensitive.

The optimal price under binding Age-Bands is similar to that under Age-Pooling, except the markup term is now inversely related to weighted demand $\bar{S}$, and the weight on each cost term is given by θ. The first-order condition thus takes into account that the price for the high-cost group is θ times that for the low-cost group. If the low-cost group (young) is more price sensitive than the high-cost group (old), there are two reasons for the high-cost group to prefer a pooling or pseudo-pooling arrangement. First, as always, more low-risk types have lower costs. However, more price-sensitive individuals also lower the optimal markup of the insurer. We use these first-order conditions for price setting in the counterfactual exercise that follows.

We then adapt this pricing rule to account for the multiple products firms offer, as firms take into account changing price on one plan has on their other plans' enrollment. Let firms offer N plans, and let $\vec{p}_i, \vec{c}_i, \vec{s}_i$ be the $1 \times N$ vectors of prices, costs, and enrollment for age group i (=L or H). Denoting the $N \times N$ matrix of cross price derivatives for consumers of type i as $\mathbf{M}_i$, the optimal price for the young group for any θ is given by:

$$\begin{aligned} \vec{p}_L^{Band} &= \left(\theta^2 \mathbf{M}_H + \mathbf{M}_L\right)^{-1} \left(\theta \mathbf{M}_H \vec{c}_H + \mathbf{M}_L \vec{c}_L\right) \\ &\quad - \left(\theta^2 \mathbf{M}_H + \mathbf{M}_L\right)^{-1} \left(\vec{s}_L + \theta \vec{s}_H\right). \end{aligned}$$

The first-order condition implies that the prices are related to both the mix of cost-type and preferences across different groups of consumers.

5 The Effect of Alternative Regulations

5.1 Simulation Method and Assumptions

In this section, we examine how alternative age-based pricing regulations would affect prices and welfare on the Massachusetts HIE. We first examine how age-based pricing interacts with imperfect competition, and compare the predicted prices to those we would expect under perfect competition. Then, we examine how changes in the maximum allowable price ratio (θ) affects transfers between the young and old, as well as between consumers and firms. We then show how minimum loss ratio regulation would interact with restrictions on age-based pricing.

Simulating firm prices requires that we specify firms' 1) demand curve, 2) costs, and 3) pricing rule. Consumer demand is derived from the preferences estimated in Table 2.[21] We use the multiproduct pricing rule developed in the previous section, and simultaneously solve the system of optimal pricing equations for a new vector of prices under each set of regulations. We infer costs from the Medical Expenditure Panel Survey, as described below. (An alternative method to infer insurer costs from observed prices under the assumption of profit maximization. That route is problematic for two reasons. However, the prices we observe are constrained by age-pricing regulation, limiting our ability to identify the relative costs of young and old. Moreover, since insurers are using coarse pricing rules, the assumption of perfect profit maximization is problematic.)[22]

Insurer costs are an important component of our simulations, but they are not directly observed. To estimate insurer costs, we rely on data from the 2008 Medical Expenditure Panel Survey (MEPS) on the health costs of different groups. To construct the table, we restrict the sample to individuals 27-64 and with moderate to high incomes and private insurance, to mimic the population in the Massachusetts Commonwealth Choice program. In the MEPS data, older consumers have higher medical expenditures, but also pay a higher percentage of those medical expenditures out of pocket. Therefore, as a measure of relative costs *to the insurer*, we form the ratio of insured costs of older groups to the insured costs of the average insured costs of 27 to 30 year-old consumers.[23] For our simulations, we assume that insurer's cost for a bronze plans is 60% of the total age-specific health costs, 70% for silver and 90% for gold plans; these correspond to the approximate actuarial values of each type of plan as set by the Connector.

The ratio of insured expenditure for the oldest consumer group (55-64 year olds) relative to those 30 and under is 2.7, implying that insurers would be constrained by a θ of 2 even in the absence of price discrimination motives. Yet, the cost ratios for slightly younger consumers (i.e. 50-54 or 45-49) is much lower (about 1.5). This suggests that price discrim-

[21] For simplicity, this specification does not allow for any unobserved heterogeneity, though the results are robust to alternative specifications.

[22] However, if the researcher is willing to assume a ratio of costs between the oldest and youngest consumers, costs can then be estimated. If we assume a cost ratio of 2.75:1, the costs estimates and counterfactuals we run are very similar to those we get when we use the MEPS. Results are available in the Online Appendix.

[23] A limitation of this analysis is that it does not account for differential selection into the exchange: the consumers who lacked coverage in the employer-based market are not representative of the population. However, in the absence of better cost data, it provides a useful baseline.

ination explains part of the pricing pattern in the data. Specifically, consider 45-49 year old consumers. Cost estimates indicate that these consumers cost only slightly more (20%) to the insurer than consumers 27-30, yet premiums are 40% higher. This is easily rationalized by differences in elasticities: consumers age 27-30 have an elasticity that is over twice the elasticity for the older group.

5.2 Unconstrained Prices with Perfect and Imperfect Competition

Figure 3 shows our estimate of what prices would be if firms did not face age-based pricing regulation. It plots premiums by age under both perfect and imperfect competition, assuming the costs in the MEPS data, the preferences in the demand system, and no regulation. Differences in preferences amplify this difference in costs, leading older consumers to have a much larger gap between prices under perfect and imperfect competition. The simulated prices under imperfect competition are more extreme than actually observed in the market (the prices for older consumers are higher than observed, and the prices for younger consumers are lower than observed), because the prices are unconstrained by age-based pricing regulation in this simulation. Note that margins on the oldest consumers can be quite large at around \$100 per month, or 20% of the purchase price. By contrast, the margins on the youngest consumers are quite slim.

5.3 Effects of Alternative Age-Bands

Here, we examine the effect of alternative age-bands (value of θ), and consider maximum allowable price ratios that range from full pooling ($\theta = 1$) to the ACA's maximum value ($\theta = 3$).To conduct our simulation of alternative age-bands, we split the sample at age 45 into young and older consumers, and use the parameter γ to indicate the relative cost of older consumers, such that $c_H = \gamma c_L$. Following the MEPS we set $\gamma = 2.75$, with average total medical expenditures of \$2500/year for younger consumers.

Table 5 describes the simulation results, presenting change in premiums relative to the baseline of Age-Unconstrained prices. We find that conditional on the age-band being set at the cost ratio ($\theta = \gamma = 2.75$), the prices for the younger group are 7% higher than they would be under Age-Unconstrained pricing. The Massachusetts age-band of 2 leads to 22% higher premiums for younger consumers conditional on a cost ratio of 2.75. We conduct additional simulations to isolate the effect of preference heterogeneity, in which we assume a maximum price ratio θ and a cost ratio γ both at 2. In that case, the prices for the young are still 10% higher than they would be under Age-Unconstrained pricing. Thus, for reasonable parameters, it seems that the transfers from younger to older consumers are due in equal parts to cost differences and preference differences.

Changes in age-based pricing do not merely transfers between consumers–they alter total consumer surplus as well. Table 5 compares consumer surplus under bands of 1 (full pooling), 2, 2.75, and 3 to unregulated pricing to capture the welfare impact of regulation. In each case, consumer welfare is higher with the pricing regulation in place than it would be in the absence of regulation; the positive compensating variation for the older consumers is larger in magnitude that the negative compensating variation for younger consumers, who must be paid to be made whole.

Insurers price to the marginal, rather then average consumer. Since the marginal consumer is likely to be a young consumer, prices fall when age bands are narrower than the cost differentials. Recall the equation for prices under Age-Bands in Proposition 1. The cost term represents the (weighted) cost of the marginal consumer, and is less than a simple average cost whenever the ratio of the price sensitivity of the younger, more price sensitive group to the older, less price sensitive group is sufficiently large.[24]

Premiums for the older consumers fall more than the premiums of the younger consumers rise (even though the older consumers value the reductions less). In our simulations, the unconstrained average prices are around \$200/month for the younger consumers and \$600/month for older consumers. If you were simply to take an average in the Age-Pooling case, you would expect a \$200/month increase in the premiums of younger consumers. But because insurers price to the marginal consumer, the increase in premiums on younger consumers is only \$100/month, for an average premium of only \$300/month in the full pooling case.

The regulation affects firm profits as well. The bottom row of Table 5 shows that firm profits fall (compared to unconstrained pricing) as the maximum allowable price ratio is lowered. The 2:1 age bands allow consumers to capture more of the surplus generated without leading to negative profits on the part of insurers. Total surplus (consumer surplus plus firm profits) increases as the maximum allowable price ratio is lowered because the price vector in equilibrium more closely reflects prices based on the cost of the marginal consumer. This allows consumers to sort more efficiently across plans.[25]

[24]Specifically, larger than $\theta - \theta^2$, which is true in all our simulations in which the maximum allowable price ratio binds. In fact, in the community rating case, the insurers would lose money. This provides a rationale for why community rating arrangements that would be possible under perfect competition may be more likely to fail under imperfect competition.

[25]Section 6 also shows age-based pricing regulation can have an effect on efficiency when the mandate is dropped or ineffective.

5.4 Minimum Loss Ratio Regulation

Minimum loss ratios (MLR) require that insurers pay out a certain percentage of premiums in health costs or refund the difference to consumers, creating a pseudo-price cap. While the Massachusetts insurers were not subject to minimum loss ratios during our time period, we simulate the effect of the ACA's 80% MLR and how it interacts with age-based pricing regulation.[26]

Table 6 shows how premiums change for younger consumers with and without MLRs under different levels of the maximum allowable price ratio. The comparison case is a minimally regulated market, with no restrictions on age-based pricing, no MLR, no subsidies; the table shows zero change in that case.

MLRs mitigate the transfers from younger consumers to older consumers that would otherwise occur under binding age-based pricing regulation. The bottom row shows the current Massachusetts age-band of 2: in the absence of MLR regulation, there are price increases of approximately 22% relative to the base case, but the introduction of the MLR regulation lowers that increase to about 17%. Intuitively, the MLR regulation dampens the incentive for insurers to set higher prices for younger consumers to gain slack against the age rating regulation. In the face of the MLRs, the insurers cannot collect the higher margins on older consumers.

6 Market Participation Without a Mandate

6.1 Participation Model

Removing the mandate to purchase insurance and allowing for non-participation in the market increases the impact of modified community rating regulations. Thus far, we have assumed all consumers purchase insurance: that the mandate is effective. In the absence of a mandate, consumers may opt out of coverage. Indeed, just as the market can unravel when costs differ, the market can also unravel when preferences differ.

In this simulation, we show how pricing under age-bands can lead to selection into and out of the exchange. The participation rate affects both market shares, which affect optimal markups, and costs. When limited by maximum allowable price ratios, insurers raise prices

[26] We impose MLRs at the insurer-state level. This level of aggregation is likely to be used by states to reduce noise and administrative complexity. MLR regulations would have slightly different effects if they instead were implemented plan-by-plan or age-by-plan group. We mechanically calculate the loss ratio using an enrollment weighted average across plans, and impose the restriction that insurers cannot set prices such that they violate the MLRs, abstracting from the possibility that the firms may end up issuing refunds. In addition, we abstract from the idea that firms trade-off Type I and Type II error when setting prices given that they want to get as close as possible to the MLR without violating the regulation.

on younger consumers. Some younger consumers opt out of the market completely, leading prices to adjust and exacerbating the transfers from the younger consumers left in the market to older consumers.

We expand the model to allow consumers to opt out of the market. Participation rates depend on a group's take-up elasticity and the prices they face. Denote the participation rate of consumer group i under an age band of θ by $\rho_{\theta i}$, and their take-up elasticity by ε_i. Further, let p_i^0 represents the price (or price index) under unconstrained pricing and $p_{i\theta}^*$ represents the optimal price (or price index) for consumer group i under a MAPR of θ. We assume full participation under unconstrained pricing as a benchmark. The participation rate can then be written as

$$\rho_{i\theta} = 1 - \varepsilon_i \left(p_i^0 - p_{i\theta}^* \right)$$

We allow participation to vary by health status and allow health status to vary by age. We split the sample by age again as young or old (under- versus over-45). We then split by good or bad health status as taken from the MEPS, and group individuals in "excellent", "very good", or "good" health together as one group, and those in "fair" or "poor" health as another. There are thus four demographic groups, each an age-health-status cell.

We need a number of additional assumptions for this simulation as well. First, we take extensive margin elasticities for both relatively healthy and relatively sick (in "fair" or "poor" self-reported health) from the CBO (2005): they are -.57 and -.34, respectively. In addition, we assume that younger consumers in "fair" or "poor" health are 3.75 times more costly to insure than consumers in "excellent", "very good", or "good" health, again following MEPS estimates.[27] We continue to assume that average cost of older v. younger individuals are given as in the MEPS.

6.2 Simulation with Mandate Removed

Table 7 describes results of removing the mandate and allowing non-participation. Selection into the exchange leads to prices that are dramatically higher, regardless of the age-band. In addition, we see that healthy, young consumers are likely to leave the market. Under the Age-Pooling regulation, approximately 60% of healthy, young consumers opt out of the market, along with 45% of the sicker consumers. In addition, there is significant unravelling with a maximum allowable price ratio of 2.75, and none of the transfers are due to differences in costs alone.

While full unravelling does not occur given these parameter estimates, consumers are

[27] The relative costs of the healthy older individuals versus sick older individuals does not play a role in our simulations, as the older individuals continue to participate in the market.

made worse off by the higher prices and there is an inefficiency resulting from uninsurance. Hence, without a mandate, there is a trade-off between the welfare benefits of pooling (as in the previous section) and the costs of uninsurance. Finally, we note that a statutory penalty can reduce or eliminate opting out of insurance. For example, a statutory penalty equal to half of the cheapest bronze premium (as it is set in Massachusetts) leads to no uninsurance in the simulation above.

The optimal price tends to increase when the mandate is removed for three reasons. First, in our simulation, the consumers that leave the market are younger consumers, who are both less expensive to insure and more price-sensitive: prices for older consumers are always lower under age-based pricing regulation than with unconstrained pricing. In addition, with the larger exit of healthy consumers from the market, average cost to insure goes up even further. Finally, the marginal consumer changes. Since there are simply fewer younger consumers in the market, the marginal consumer is more likely to be older, and more expensive to insure. This further increases prices.

6.3 Death Spiral Even With Perfect Risk Adjustment

The HIE market can unravel due to differences in preferences alone if the participation elasticity is high enough. To show this, we conduct a simulation of what can occur even with perfect risk adjustment, in which insurers' costs do not vary by age of enrollee.[28] As a result, prices would not vary by age due to costs, but they will vary due to a price discrimination rationale. Under Age-Pooling regulation without an effective mandate, the young would face higher markups, and some would substitute out of the market.

We simulate the effect of allowing for opting out in the absence of a mandate in Figure 4. Using intermediate values from the previous simulations, we assume firms' optimal markup is 15% for the younger consumers, and 35% for the older consumers. We assume full market participation among the older group (they always face lower prices than they would if insurers were unconstrained), but consider a range of participation elasticities for the young consumers. Estimates of health insurance take-up elasticities in the literature vary substantially, from near zero to -2 (Washington and Gruber 2005, Cutler and Reber 2002). However, elasticities from the employer-sponsored insurance may not correspond well to this new environment. Therefore, we simulate the optimal markup under take-up elasticities ranging from zero (full participation) to -5. (We use optimal markups for each group as estimated above, and assume equal population shares for the young and old.)

We can get a "death spiral" from differences in preferences alone for a take-up elasticity

[28] Note, there is no risk adjustment in the Massachusetts HIE during our sample.

of -5 and larger:[29] no younger consumers participate in the market and the optimal pooled markup is equal to the optimal markup for older consumers. Our results emphasize the heterogeneity in consumer preferences for insurance that has been noted in the literature, and connects this idea to the response of insurers facing regulation in a new market, a health insurance exchange. Modified community rating rules have a large impact on this market, even in the presence of a mandate. However, if the mandate were dropped or were not effective, heterogeneity in preferences alone can lead to a death spiral effect in which all price-sensitive consumers exit the market.

7 Conclusion

This paper has analyzed consumer behavior and pricing regulation using a novel data set in a health insurance exchange that serves as a model for national health reform. We find strong evidence of heterogeneity in preferences, with younger consumers who are more than twice as price sensitive than their older counterparts. Price discrimination gives insurers a motive, in addition to costs, to increase premiums on older consumers, amplifying variation in insurance prices.

Our simulations show that age-based pricing restrictions increase consumer surplus by leading insurers to price to marginal, young, inexpensive consumers, but transfers money from younger consumers to older consumers. However, minimum loss ratios can mitigate these transfers while retaining the incentive for insurers to price to the marginal, rather than average consumer, when setting prices. With the ACA's maximum allowable price ratio and minimum loss ratios, consumer surplus is increased by relative to unconstrained age pricing. However, partial unravelling and dramatic premium increases can occur in the absence of an effective mandate. This furthers the importance of statutory penalties, which may mitigate these concerns.

For each of the regulations we consider–modified community rating and age-based pricing limitations, minimum loss ratios, and mandates–assuming a perfectly competitive market would have led to misleading results. Insurance market regulations must not only consider the nature of consumer demand, but also strategic insurer pricing in the face of consumer demand. Our results from Massachusetts provide a starting point for researchers who study HIE and related markets, as well as for regulators designing the exchanges.

[29] So long as the older consumers are less responsive to price than the younger consumers, a death spiral will exist if the participation elasticities of the younger group are large enough.

References

Blumberg, L., Buettgens, M., and Garrett, B. (2009). Age rating under comprehensive health care reform: Implications for coverage, costs, and household financial burdens. *Timely Analysis of Immediate Health Policy Issues. Available at: http://www. urban. org/UploadedPDF/411970_ age _rating. pdf.*

Carlin, C. and Town, R. (2007). Adverse selection, welfare and optimal pricing of employer-sponsored health plans. *U. Minnesota Working Paper*.

Chu, C., Leslie, P., and Sorensen, A. (2011). Bundle-Size Pricing as an Approximation to Mixed Bundling. *The American Economic Review*, 101(1):263–303.

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

Dafny, L. (2010). Are health insurance markets competitive? *American Economic Review*, 100(4):1399–1431.

Dafny, L., Duggan, M., and Ramanarayanan, S. (forthcoming). Paying a premium on your premium: Consolidation in the u.s. health insurance industry. *American Economic Review*.

Dafny, L., Ho, K., and Varela, M. (2010). Let them have choice: Gains from shifting away from employer-sponsored health insurance and toward an individual exchange. *National Bureau of Economic Research Working Paper 15687*.

Dafny, L. and Ramanarayanan, S. (2012). Does it matter if your health insurer is for-profit? effects of ownership on premiums, insurance coverage, and medical spending. *National Bureau of Economic Research Working Paper 18286*.

Einav, L., Finkelstein, A., and Cullen, M. (2010). Estimating welfare in insurance markets using variation in prices. *Quarterly Journal of Economics*, 125:877–921.

Ericson, K. M. (2012). Consumer inertia and firm pricing in the medicare part d prescription drug insurance exchange. *National Bureau of Economic Research Working Paper 18359*.

Ericson, K. M. and Starc, A. (2012). Heuristics and heterogeneity in health insurance exchanges: Evidence from the massachusetts connector. *The American Economic Review*, 102(3):493–497.

Finkelstein, A. (2004). Minimum standards, insurance regulation and adverse selection: evidence from the medigap market. *Journal of Public Economics*, 88(12):2515–2547.

Finkelstein, A., Poterba, J., and Rothschild, C. (2009). Redistribution by insurance market regulation: Analyzing a ban on gender-based retirement annuities. *Journal of financial economics*, 91(1):38–58.

Frank, R. and Lamiraud, K. (2009). Choice, price competition and complexity in markets for health insurance. *Journal of Economic Behavior & Organization*, 71(2):550–562.

Geruso, M. (2011). Community rating in employer health insurance: Inefficiencies beyond adverse selection. *Working Paper*.

Gruber, J. (2010). The tax exclusion for employer-sponsored health insurance. *National Bureau of Economic Research Working Paper 15766*.

Gruber, J. (2011). Massachusetts points the way to successful health care reform. *Journal of Policy Analysis and Management*.

Gruber, J. and Washington, E. (2005). Subsidies to employee health insurance premiums and the health insurance market. *Journal of Health Economics*, 24(2):253–276.

Handel, B. (2011). Adverse selection and switching costs in health insurance markets: When nudging hurts. *National Bureau of Economic Research Working Paper 17459*.

Health Connector (2010). Report to the Massachusetts Legislature: Implementation of Health Care Reform, Fiscal Year 2010.

Ilayperuma Simon, K. (2005). Adverse selection in health insurance markets? evidence from state small-group health insurance reforms. *Journal of Public Economics*, 89(9):1865–1877.

Ketcham, J., Lucarelli, C., Miravete, E., and Roebuck, M. (2010). Sinking, swimming, or learning to swim in medicare part d. *American Economic Review*.

Kolstad, J. and Kowalski, A. (2010). The impact of an individual health insurance mandate on hospital and preventive care: Evidence from Massachusetts. *National Bureau of Economic Research*.

Long, S. and Phadera, L. (2009). *Estimates of health insurance coverage in Massachusetts from the 2009 Massachusetts Health Insurance Survey*. Commonwealth of Massachusetts, Division of Health Care Finance and Policy.

Maestas, N., Schroeder, M., and Goldman, D. (2009). Price variation in markets with homogeneous goods: The case of medigap.

Nevo, A. (2000). A practitioner's guide to estimation of random-coefficients logit models of demand. *Journal of Economics & Management Strategy*, 9(4):513–548.

Sethi-Iyengar, S., Huberman, G., and Jiang, W. (2004). How much choice is too much? Contributions to 401 (k) retirement plans. *Pension design and structure: New lessons from behavioral finance*, pages 83–95.

Simon, K. (2005). Adverse selection in health insurance markets? evidence from state small-group health insurance reforms. *Journal of Public Economics*, 89:1865–1877.

Starc, A. (2010). Insurer pricing and consumer welfare: Evidence from medigap. *Manuscript. Harvard University*.

Strombom, B., Buchmueller, T., and Feldstein, P. (2002). Switching costs, price sensitivity and health plan choice. *Journal of Health economics*, 21(1):89–116.

Zuckerman, S. and Rajan, S. (1999). An alternative approach to measuring the effects of insurance market reforms. *Inquiry*, pages 44–56.

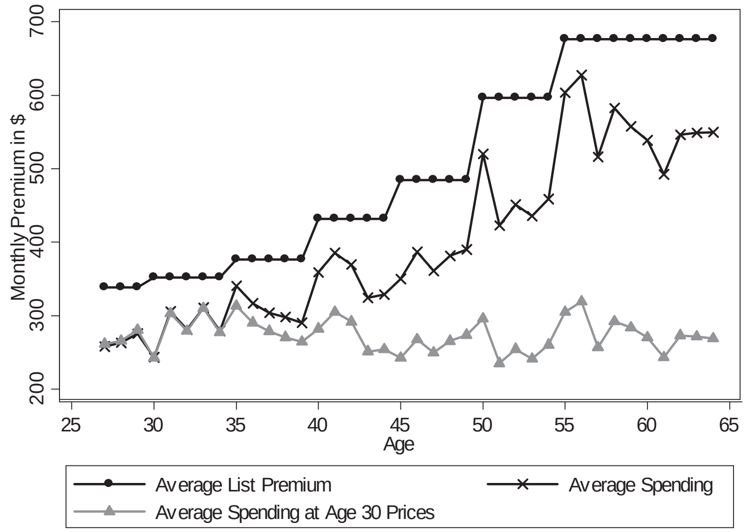

Figure 1: Average Prices and Spending, By Age. Notes: Data: Nov. and Dec. 2009. Average list premium is the plan-weighted average using the Nov. plan-zipcode price. Average spending is person-weighted. Average spending at age 30 prices uses actual choices but prices for a 30-year old in Nov. in that zipcode.

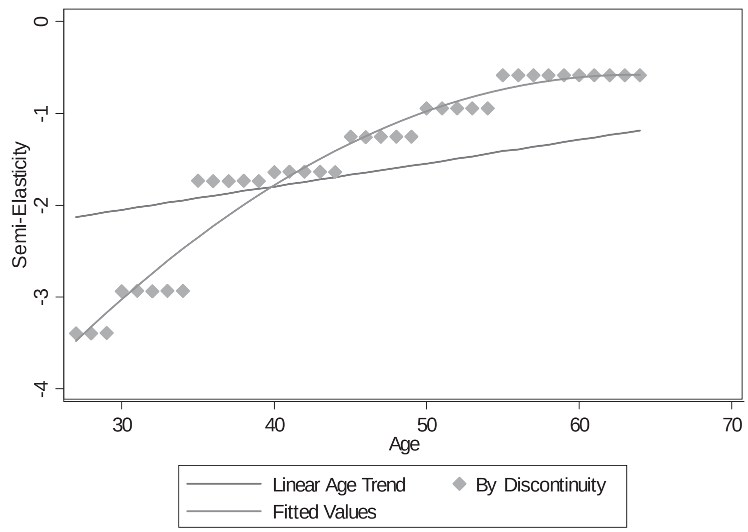

Figure 2: Semi-Elasticities By Age. Notes: Semi-elasticity gives the percent reduction in market share resulting from a $100 increase in monthly premium. Linear age trend is plotted using results from Column 3 of Panel A of Table 3. Discontinuity results plot the average semi-elasticity obtained from Panel B of Table 3. The fitted values fit a quadratic trend to these estimates.

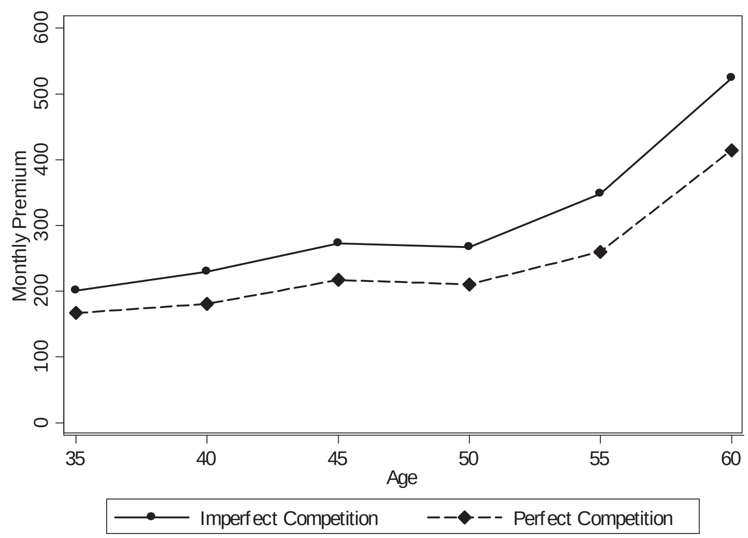

Figure 3: Simulated Monthly Premiums With Unconstrained Pricing. Notes: Weighted average based on observed choices. Assumes no age-based pricing regulation. Under perfect competition, we assume insurers charge at cost, with costs taken from the 2008 MEPS (see Table 6). Under imperfect competition, we assume firms charge the optimal markup based on age-specific price sensitivity calculated in Table 3 Panel B.

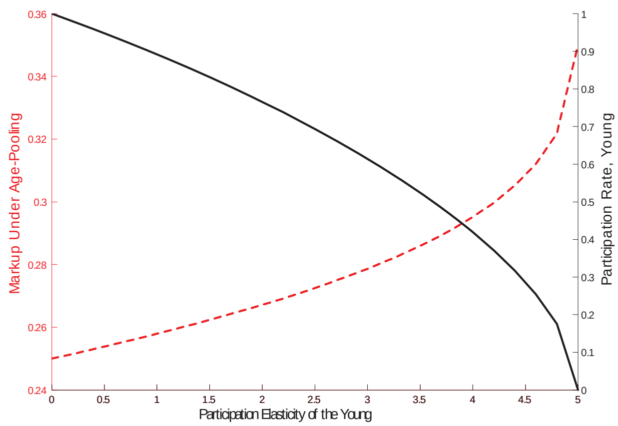

Figure 4: Optimal Markups and Non-Participation in the Absence of a Mandate.

Table 1: Demographics of the Connector

	Full Sample	Nov-Dec 2009
	Demographics	
Age	35.93	36.15
% Female	0.4524	0.4925
# of Lives Covered	1.311	1.291
Premium Paid (Monthly)	349.77	375.3
	Tiers	
Bronze	39.65	37.87
Bronze Plus	1.8	3.09
Silver	13.24	14.01
Silver Plus	2.86	2.36
Silver Select	5.16	5.51
Gold	7.42	7.01
Young Adult	29.87	30.15
	Insurers	
Blue Cross Blue Shield of Massachusetts	31.47	30.55
Fallon Community Health Plan	15.38	17.94
Harvard Pilgrim Health Care	22.7	21.19
Health New England	2.76	2.84
Neighborhood Health Plan	20.51	20.33
Tufts Health Plan	7.18	7.15

Note: Numbers represent simple averages from the raw data.

Table 2: Price Sensitivity by Age in Conditional Logit Model

	Panel A: Basic Conditional Logits (All Ages)					
	(1)		(2)		(3)	
Premium	-0.357***		-2.018***		-2.266***	
(in $100s)	(0.122)		(0.306)		(0.369)	
Premium*age			0.0298***		0.0267**	
			(0.00488)		(0.0114)	
Fixed Effects	Plan		Plan		Plan, Plan*Age	
N Person*Plan	20,838		20,838		20,838	
	Panel B: Conditional Logits by Age Group					
	27-34	30-39	36-44	40-49	46-54	50+
Premium	-3.574***	-2.611***	-2.354***	-2.271***	-1.512***	-1.234***
(in $100s)	(0.533)	(0.560)	(0.606)	(0.508)	(0.572)	(0.316)
N Person*Plan	8,512	5,396	4,380	4,459	3,745	5,628
Plan and Tier*Age^2	Yes	Yes	Yes	Yes	Yes	Yes
	Panel C: Conditional Logits For Counterfactual Exercise					
		Under Age 45			Age 45+	
Premium		-2.747***			-0.752***	
(in $100s)		(0.382)			(0.266)	
Plan and Tier*Age^2		Yes			Yes	
N Person*Plan		12,892			7,946	

Note: Heteroskedasticity robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Panels B and C include plan fixed effects, and tier effects interacted with age trends (both linear and quadratic terms).

Table 3: Age-based Price Sensitivity in Additional Specifications

	(1)	(2)	(3)	(4)
	CL	CL	NL	ML
Premium in Hundreds of $	-3.211***	-3.207***	-1.398***	
	(0.416)	(0.456)	(0.165)	
...*1(30-34)	0.602***	0.440*	0.138	0.475***
	(0.181)	(0.237)	(0.119)	(0.180)
...*1(35-39)	0.779***	0.461	0.195	0.585***
	(0.255)	(0.378)	(0.130)	(0.202)
...*1(40-44)	0.938***	0.572	0.326***	0.609***
	(0.328)	(0.461)	(0.121)	(0.210)
...*1(45-49)	1.084***	0.780	0.218*	0.470**
	(0.390)	(0.492)	(0.121)	(0.216)
...*1(50-54)	1.471***	1.305**	0.398***	0.736***
	(0.450)	(0.515)	(0.128)	(0.242)
...*1(55+)	1.892***	1.855***	0.707***	1.197***
	(0.500)	(0.519)	(0.0994)	(0.211)
Bronze Dissimilarity Parameter			0.531***	
			(0.0777)	
Silver Dissimilarity Parameter			0.608***	
			(0.0933)	
Gold Dissimilarity Parameter			0.977***	
			(0.208)	
Mean Premium Parameter				1.019***
				(0.0945)
S.D. Premium Parameter				0.345***
				(0.0400)
Fixed Effects	Plan	Plan		Tier
	Plan*Age	Plan*Age	Tier*Age	Tier*Age
		Plan*Age^2	Insurer	Insurer
N Person*Plan	20,838	20,838	20,838	20,838

Note: Heteroskedasticity robust standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1. Columns 1 and 2 contain conditional logit specifications. Column 3 estimates and nested logit and reports dissimilarity parameters, where a dissimilarity parameter equal to one collapses to a conditional logit. Column 4 reports a mixed logit specification in which the price coefficient is allowed to take on a lognormal distribution. Additional mixed logit specifications by age confirm the general pattern in Table 3.

Table 4: Comparison of Costs Across Age Groups

	Population (in thousands)	Avg. Annual Expense ($ per capita)	Percent Paid... Out of pocket	by Private Insurance	Avg. Insurer Cost c_a	Ratio: c_a/c_{27}
Total	98968	$ 3,992	18.5	76.5	$ 3,054	
Age group:						
27-30	7226	$ 2,401	17.4	79.7	$ 1,914	1.00
31-34	11715	$ 2,509	18.3	77.5	$ 1,945	1.02
35-39	12866	$ 2,723	23.3	73.9	$ 2,012	1.05
40-44	13863	$ 3,279	18.7	76.6	$ 2,512	1.31
45-49	13652	$ 3,241	19.6	71.6	$ 2,320	1.21
50-54	15092	$ 4,046	18.9	75.9	$ 3,071	1.60
55-64	24554	$ 6,627	17.2	78.1	$ 5,175	2.70

Note: Data taken from 2008 MEPS, with authors' calculations. Sample selection: people age 27-64 with middle or high incomes with any private insurance. Avg. insurer cost is mean private insurer expenditure for this sample.

Table 5: Surplus Change, Relative to No MAPR Regulation

	Max. Allowable Price Ratio			
Compared to Unconstrained Price:	1	2	2.75	3
% Change in Premium:				
...Older Consumers	-49.92%	-16.38%	-3.75%	-1.01%
...Younger Consumers	33.22%	22.31%	6.78%	1.65%
Surplus Change, $ per consumer:				
...Older Consumers	298.37	93.99	21.18	6.33
...Younger Consumers	-49.70	-37.36	-12.16	-3.25
...Total	248.67	56.63	9.02	3.08
Change in Firms' Profits ($ Per Consumer)	-124.33	-28.31	-4.51	-1.53

Note: Source: Data and authors' calculations. Taken from a series of counterfactuals in which the maximum allowable price regulation (MAPR) is altered, as described in the text. Surplus change represents the compensating variation required to provide consumers with the same level of utility relative to unconstrained premiums. Percentage increase or decrease in each age group's premium is taken as a percent of their premium under $\theta = \infty$.

Table 6: Percent Change in Premiums for Younger Group, Relative to a Minimally Regulated Market: Results of Simulations

	Minimum Loss Ratio	
	Yes	No
MAPR = ∞	-9.1	0.00
Cost MAPR (2.75)	-0.50	6.8
Mass. MAPR (2)	16.5	22.3

Source: Data and authors' calculations. Taken from a series of counterfactuals in which the impact of various regulations is examined simultaneously, as described in the text. "MAPR" represents the maximum allowable price regulation. All percentages reported are the change in price for the younger group of consumers.

Table 7: Participation in the HIE without a Mandate: Results of Simulations

MAPR	With Mandate: % Increase in price	Without Mandate: % Increase in price	% of Healthy Young Participating	% of Sick Young Participating
1	33.2	107.8	40.1	64.3
2	22.3	69.0	59.5	75.9
2.75	6.8	39.1	79.6	87.8
3	1.7	30.4	85.5	91.4

Source: Data and authors' calculations. Taken from a counterfactual in which healthy and sick younger consumers opt out of the market at different rates, as described in the text. The first two columns represent the percent increase in price that results in moving from unconstrained pricing to full pooling. The increase is larger in the absence of the mandate because healthier consumers opt out of the market. The second two columns represent the percentage of consumers (of each type) participating in the market.