NBER WORKING PAPER SERIES

AGE-BASED HETEROGENEITY AND PRICING REGULATION 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

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.

Age-Based Heterogeneity and Pricing Regulation on the Massachusetts Health Insurance Exchange
Keith M. Marzilli Ericson and Amanda Starc
NBER Working Paper No. 18089
May 2012
JEL No. I11,I13

ABSTRACT

Little is known about consumer behavior or insurer incentives in health insurance exchanges. We analyze choice on the Massachusetts exchange, using coarse insurer pricing strategies to identify price sensitivity. We find substantial age-based heterogeneity: younger individuals are more than twice as price sensitive as older individuals. Modified community rating regulations interact with price discrimination, as our results imply higher markups on older consumers. Age-based pricing regulations would bind even conditional on perfect risk adjustment, highlighting the importance of considering insurer incentives when regulating insurance markets. Changes in age-based pricing regulation can result in transfers of 8% of the purchase price.

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 open marketplaces for private insurance - raise new questions about the effects of regulation in insurance markets. They also provide new opportunities to study consumer demand in a context with a wide range of choice. Traditionally, most individuals received either employer-based health insurance or government-provided health insurance (Medicare or Medicaid). HIEs are changing the way people purchase health insurance in that they combine retail and regulatory functions. States may set up exchanges to comply with the 2010 Patient Protection and Affordable Care Act (PPACA)[1] or as a result of their own reforms (Massachusetts and Utah). These states will have substantial latitude in designing and regulating these exchanges, and they will make choices that will shape the market for individually-purchased health insurance. 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.

This paper examines consumer demand on the Massachusetts HIE, known as the Connector. The Connector provides an early look at a comprehensive HIE in action, as it was the first HIE established in the United States and has been providing coverage since 2007. Other markets also offer insight into HIEs, but have crucial differences. While Medicare Part D is like an insurance exchange, it offers a limited type of coverage (prescription drug) to a narrow age range (the elderly).[2] Employer-sponsored insurance at large employers may offer a range of choice akin to an HIE, yet it differs in regulation (e.g. plans cannot price differentially by age) and in the nature of competition (the employer negotiates directly with insurers). We add to a growing literature on choice in these contexts. Dafny, Ho, and Varela (2010) examine existing

[1] The 2010 Patient Protection and Affordable Care Act (PPACA) both mandates that all Americans carry health insurance and requires that states establish HIEs to facilitate individual purchase of health insurance. If a state does not establish an HIE, consumers will be eligible to purchase via a federal HIE.

[2] See Duggan, Healy and Scott Morton (2008) on the Medicare Part D reform, and Ericson (2012) on how consumer inertia affects the design of the Medicare Part D exchange.

employer-sponsored insurance and argue that increasing choice would lead to consumer welfare gains. In the Medicare Part D market, Abaluck and Gruber (2011) find evidence of choice inconsistencies and conclude that limiting the choice set of consumers may lead to better outcomes, while Ketcham et al. (forthcoming) argue that consumers learn over time and choose better plans.

Our analysis shows that regulation of exchanges will play a critical role in insurance markets in at least two important ways. First, regulating what counts as sufficiently generous insurance to satisfy mandates for coverage - minimum creditable coverage - will determine the plan that many people get. A majority of our sample chooses the least generous coverage sufficient to satisfy the mandate, which is less generous than coverage in typical employer plans (see also Ericson and Starc 2012). Second, regulating who pools with whom through modified community rating alters the division of surplus among consumers by limiting how insurers can vary prices based on consumer characteristics. These and other regulations are set not only by federal and state legislation, but also at the level of the HIE itself. In Massachusetts, the Connector has the power to standardize plan features and determine whether a plan provides sufficient coverage to satisfy the mandate. Exchanges also control how information about insurance plans is presented to consumers, the defaults individuals face, and the frequency of the open enrollment period.

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 within limits, attempts to mitigate some of the welfare loss. A critical feature of modified community rating in Massachusetts has been age-based pricing: older individuals pay higher prices than younger individuals, but regulation limits the extent to which prices can vary by age. A large literature has analyzed the impact and importance of community rating laws, often considering the redistribution of risk and assuming markets are perfectly competitive.[3]

[3] For additional information, see Simon (2005), Zuckerman and Rajan (1999), Herring

Meanwhile, a recent literature (see Dafny et al. 2010, Starc 2010, and Lustig 2010 for examples) has highlighted insurer market power as a source of increasing health insurance rates. This paper knits together these literatures. Our analysis of consumer choice on the Connector provides the basis for examining how modified community rating regulation interacts with insurer incentives under imperfect competition. We emphasize the importance of accounting for both consumer behavior and strategic firm behavior when designing HIEs.

Identifying price elasticities is necessary to determine the potential for insurers to charge markups above cost (relevant for antitrust regulation), as well as the effect of subsidies for more generous insurance coverage (relevant for policy debates regarding the effect of the tax subsidy for health insurance). Existing estimates of price sensitivity for health insurance vary substantially by context, and most examine employer-sponsored health insurance.[4] The only work addressing HIEs specifically is Ericson and Starc (2012), whose measures of price sensitivity do not account for endogeneity of premiums; instead, that paper focuses on the role of heuristics in choice.

We analyze age-based pricing regulation on HIEs and show that prices by age vary not only because costs vary by age, but also because price sensitivity varies by age. We focus on the impact of this particular policy, as it is important in its own right, but also provide a general framework for modeling modified community rating in the presence of imperfect competition. Using both a reduced-form and discrete choice framework, we measure how sensitive consumers are to price. We use discontinuities in the ways firms set their prices by age to identify consumers' response to age: prices are constant within five year age blocks (e.g. 30 to 34), and then jump. In our reduced-form specifications, we find that total spending rises nearly one-for-one with price increases. However, this misleadingly suggests little consumer response to price; in fact, because consumers are already clustered at the cheapest plans (25% choose the cheapest plan available to them in our sample period), they have little ability

and Pauly (2006), and Finkelstein, Poterba, and Rothschild (2009).

[4] See Carlin and Town (2007), Gruber and Washington (2006), Bundorf, Levin and Mahoney (2008) and Einav, Finkelstein and Cullen (2010).

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 derived from their underlying demand for medical services.[5] Plans with the same level of coverage are not typically perfect substitutes, as insurers have different networks of providers; consumers may also attach value to firm brand and reputation. Moreover, individuals are constrained by their choice set.

We identify substantial variation in demand elasticities: younger consumers are significantly more price sensitive than their older counterparts. Furthermore, there is heterogeneity within age groups. The estimates show that consumers in the 75th percentile of price sensitivity have elasticities four times larger than their counterparts in the 25th percentile. Heterogeneity between and within observable groups is important for evaluating potential policies, such as the minimum creditable coverage regulation and age-based pricing regulation.

Insurers may price discriminate by responding to heterogeneity in consumer demand that is correlated with observable tags. (Here we study age, but other contexts could include gender, geographic location, and race.) Using our estimates, we simulate the ability of insurers to price discriminate under various models of market structure and highlight the potential welfare effects of alternative pricing regulations. In order for insurers to price discriminate, there must be consumer heterogeneity in preferences that can be both observed and priced. We conduct a counterfactual exercise in which we simulate premiums and estimate the distributional consequences of eliminating or tightening age rating rules. Ultimately, the disparities in price sensitivity by age imply that age rating rules are one of the most important regulatory features of the exchange.[6] We estimate the effect of moving from no restrictions on age-based

[5] Moreover, consumers may be daunted by the complexity of the insurance product and difficulty forcasting their future medical expenses. In such cases, consumers may rely on heuristics. We explore this idea in further detail in Ericson and Starc (2012).

[6] Geruso (2011) considers the impact that preference heterogeneity has on welfare in a model in which insurers are perfectly competitive. Our model, by contrast, shows how imperfectly competitive insurers can amplify this effect and how differences in preferences can lead to transfers in the absence of cost differences via price discrimination.

pricing to regulations that prohibit age-based pricing. Based on insurer price discrimination alone, this change leads to transfers of 8% of the purchase price of insurance; accounting for cost differences between ages leads to even larger price increases.[7]

Finally, we allow consumers to opt out of the market to capture the importance of the individual mandate. Even in the case of full risk-adjustment, so that costs do not differ by age, differences in preferences alone can lead to this market unraveling. We allow consumers to opt out of the market completely and the resulting change in the composition of consumers in the market to affect markups. If younger consumers are allowed (in the model or by law) to opt out of coverage, modified community rating can lead them to opt out of coverage even absent cost differences. As the most price-sensitive 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 until there are only price insensitive consumers left in the market.

This paper shows how the heterogeneity in consumer preferences for insurance products noted in the literature (Cohen and Einav 2003, Cutler, Finkelstein, and McGarry 2006) interacts with public policy. We identify the regulations that are of key importance in this market and explain why these regulations are critical given consumer demand and the strategic reaction of insurers to both demand and regulation. Ultimately, choosing the set of consumers who form a risk pool is critical for determining the allocation of surplus in insurance markets. We highlight the importance of considering differences in preferences as well as differences in costs. As states decide how to define their exchanges and which segments of consumers to include - such subsidized enrollees, younger consumers eligible for catastrophic plans, or employees in small groups - understanding differences in preferences is just as critical as understanding differences in risk.

The paper proceeds as follows. The second section describes the Massa-

[7] These implications - though a transfer, rather than a welfare loss - are larger than the welfare loss from selection (see Bundorf, Levin, and Mahoney 2008 and Einav, Finkelstein, and Cullen 2010) and highlight the importance of considering the incentives of imperfectly competitive insurers when designing and regulating insurance markets.

chusetts Connector, rating regulation, and some reduced-form results. Section 3 describes reduced-form evidence of consumer spending elasticity, while Section 4 expands on this analysis to incorporate a discrete choice approach. Section 5 discusses incentives for non-uniform pricing and describes the related counterfactual exercise. Section 6 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 (for their age) 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.

Finally, the reform established an unsubsidized health insurance exchange (the Commonwealth Choice program, run by the Connector) to facilitate non-group coverage purchased directly by households and small group purchase of insurance. The Commonwealth Connector Authority operates as a quasi-public agency and has offered health insurance through the Connector since May 1, 2007 (with the mandate taking effect July 1, 2007). The Connector shapes the market for individual coverage in Massachusetts in a number of ways. It operates the exchange's website[8] and chooses which features of insurance plans are highlighted.

The Massachusetts reform has been effective at reducing the number of uninsured individuals. In 2009, 97.3% of the population was insured (Long and Phadera 2009), with increases in the insured coming from individuals purchasing insurance through the Connector, 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 Connector on the level and growth rate of premiums is a point of contention. How the Connector affects insurance prices depends on both consumer demand and market structure. By characterizing both in this paper, we provide a foundation for future analysis of the impact of the Connector.

2.2 Regulation of the Health Insurance Exchange

There are two important regulations in the market:

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 critically important in a market with a mandate.[9] While MCC may

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

[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.

not be directly under the control of an exchange regulator or designer, it will dramatically shape the market within the exchange.

Modified Community Rating: Modified community rating rules apply to pricing on the exchange. Specifically, rates for the same product have to fall within a 2:1 band across ages and geography. For a given plan, the highest quoted premium can be at most twice the lowest quoted premium. 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 rating.

Figure 2 shows that insurers are clearly constrained by regulations for age-based pricing: the average monthly premium for a 27-year-olds is just over $300, which the premiums for older consumers are just over $600. The choice of rating-bands will alter the division of surplus among young consumers, older consumers, and firms. While PPACA 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 ratio of 2.8 to 1 as a rating band (Carey and Gruber 2010).

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 initially had relatively wide latitude in designing these plans, which were grouped into tiers based on actuarial value: bronze, silver, and gold. 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. Beginning in 2010, 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. In addition to this main market, there is a separate market for young adult consumers aged 18 to 26, in which plans tend

[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.

to have more limited coverage, such as optional prescription drug coverage.

Consumers face a number of steps when purchasing insurance from the Connector. (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. Prior to the 2010 standardization of plan types, consumers needed to weigh multiple dimensions of plans (copayments, coinsurance, dental coverage), even within each tier. 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 the beginning of the Connector's existence 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 Connector's existence 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. We also examine how long consumers stay in the Connector since their initial enrollment. Figure 1 gives the Kaplan-Meier survival curve for time spent in the connector, split by tier of plan chosen. Approximately half of our observations are censored because these individuals are still enrolled in insurance when our data sample ends. The median consumer is enrolled in a Connector plan for about 13 months, and there is no spike in individuals exiting the Connector after their one-year contract is complete. It shows only small differences in enrollment duration between tiers.

This paper focuses 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 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 two subsets of the data: November-December 2009, and July-December 2009. Because we observe 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 Connector website using a Perl script. Because the plan menu is relatively constant from July through December, we are able to infer the prices individuals faced from July to October with a high degree of accuracy. The Data Appendix gives more details. The choice of sample period does not have a strong effect on the results, and we show the robustness of our results to various sample selections.

3 Consumer Response to Price: Identification Strategy and Reduced-Form Evidence

3.1 Coarse Firm Pricing

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. The results from the literature are mixed: some studies, such as Cutler and Reber (1998) find relatively high elasticities among young, healthy consumers, but other studies find that demand for health insurance is typically inelastic.[11] We use a regression 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 (within broad limits) by both zipcode and age. However, firms price more coarsely than the zipcode and age level. We first use coarse geographical pricing to define markets. Instead of varying prices for each zipcode, firms set prices for larger geographic regions that roughly correspond to hospital referral networks that may be a good proxy for underlying costs.[12] 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 zipcodes in which prices do not vary within a plan-age cell. (See Data Appendix for details.)

We use coarse age-based pricing rules to identify price sensitivity. Firms do not vary prices continuously as individuals age. While prices have discrete jumps at various ages, preferences are likely to evolve continuously across ages. The listed premium given plan has discrete changes at ages ending in 0 or 5 (30, 35, 40, etc.). Figure 2 shows jumps at each age in the average premium for a constant set of plans.[13] These jumps translate into very similar consumers

[11] See the survey evidence of Kreuger and Kuziemko (2011).

[12] These differences could also be driven by differences in medical utilization.

[13] 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

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.

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 zipcodes alone gives rise to over 40,000 potential prices, all of which must be submitted to a regulator for approval. Optimally pricing each combination and submitting it for approval would be a costly exercise for firms and could trigger adverse regulatory action. Moreover, Chu, Leslie, and Sorenson (2011) show that firms can obtain profits close to a perfectly price discriminating firm using coarse pricing rules. Finally, regulators have identified these discontinuities as a potential problem and have introduced legislation to "smooth" the relationship between premiums and age, suggesting they are not a result of shifts in utilization or preference.

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.

Appendix Table A.2 supports our identification strategy by showing that characteristics of enrollees' zipcodes 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 Connector does not change at the various age cutoffs. Fig-

bronze plan varies slightly within each age category but stays between 1.8 and 2.

ure A.3 shows the number of enrollees in each one-year age bin (we do not have exact 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.[14]

3.2 Reduced-Form Evidence on Response to Price

This section examines how choice of health insurance responds to price, using reduced form evidence on total spending. When the marginal price of more generous insurance plans decreases, consumers may chose to increase the generosity their health plan. One way to summarize the response to price is the insurance spending elasticity, which relates relate total premiums paid to the list prices individuals face. We summarize the effect of a price change on insurance spending using the following model:

$$\ln y_i = \eta \ln (p_i) + \gamma \omega_i,$$

where y_i is the total insurance premiums paid by individual i (given the actual prices), p_i is a price index for a representative bundle of plans, and ω_i is a vector of individual characteristics. The insurance spending elasticity is given by η and says that if the price index rises by 1%, the total spending rises by η%.[15] If $\eta < 1$, individuals respond to higher prices by reducing their spending on insurance, while if choice of insurance plan stayed the same, then $\eta = 1$.

In this context, the percentage price increase at each threshold varies

[14] The table presents results using the November-December Analysis Sample; the results using the July-December sample are similar and, in fact, do not contain any significant differences in zipcode characteristics at age 55.

[15] Of course, while identifying η is a valuable way of summarizing the data that can facilitate out-of-context prediction, individuals do not in fact face a continuous choice of dollars spent on health insurance; the discrete choice individuals actually face is modeled in Section 4.

among plans and insurers.[16] We therefore create a price index, in which each plan is assigned a weight. Because the plan menu varies by geographic region, we create geographic-specific weights: a plan's weight is the fraction of people in a geographic region who chose that plan, averaged over July to December 2009.[17] Column 1 of Table 2 shows how the price index jumps at each age threshold. It presents the results of the following regression:

$$\ln p_i = G(a) + \sum_{s \in 1, \ldots 7} 1_{a \geq a_s^*} \pi_s + \gamma \omega_i, \tag{1}$$

where $G(a)$ is a linear spline in age and ω_i includes gender, month of enrollment, and indicators for geographic region.[18] The coefficients π_s multiply indicator variables for whether age is greater than or equal to each of the age thresholds (each value of a_s^*) used for pricing. Each value of π_s shows how the price index jumps at the threshold a_s^*: for instance, we see that the price index increases by 20.4 log points when an individual turns 50. The jumps in prices are relatively small at age 30 and 35 but are more substantial at older ages.[19]

Next, we examine how total spending on premiums changes at each age threshold shown in Column 2 of Table 2. It presents the results of the regression

$$\ln y_i = G(a) + \sum_{s \in 1, \ldots 7} 1_{a \geq a_s^*} \kappa_s + \gamma \omega_i \tag{2}$$

where $G(a)$, ω_i, and the age category indicators are the same as in Equation 1. The values of κ_s from this regression show how spending on premiums jumps

[16] By contrast, a change in tax deduction for employer-sponsored health insurance (as in Gruber and Washington 2005) would lead to the same percentage change in price for all the plans, eliminating the need to construct a price index.

[17] To construct a reasonable price index, we exclude geographic regions that had fewer than 10 zipcodes, as well as geographic regions that had fewer than four insurers.

[18] Because we know the pricing model, gender and the linear age spline do not predict prices; they are included for comparability with later regressions.

[19] The measured increases in the price index at the age thresholds (and their relative magnitudes) vary depending on how the price index is constructed. We gave plans weights based on popularity in a geographical region. Ideally, we would assign age-group specific weights (to create a Laspeyres or Paasche index). However, the sparsity of the data makes this an unappealing route. The construction of the price index has no bearing on how we measure the change in spending at each threshold (but is relevant for estimating η).

at each age discontinuity. Thus, we see spending on premiums jump 19.2 log points at age 55, controlling for linear age trends above and below 55, along with other variables.

Comparing the percentage increase in spending (κ_s in Column 2) to the percentage increase in the price index (π_s from Column 1), we see that the increase in spending is slightly less than the increase in prices for all age thresholds except age 55. The bottom panel of Table 2 shows the results of an instrumental variable regression that instruments for price index by age-category, controlling for an age spline.[20] Column 1 is thus the first-stage of this IV regression, and the F-statistic for excluded instruments (age-discontinuities) is a substantial 2823. The resulting estimate of η is 0.962 (s.e. 0.176), indicating that a 10% increase in the price index leads to a 9.6% increase in total spending in this population - a relatively limited response by individuals.[21]

The pattern of results above is robust to larger samples. We limited our sample to the July through December 2009 sample, because outside that time period we do not have the full menu of prices an individual faced. Nonetheless, we can run an (imperfect) analog of equation 2, where the unit of observation is a plan chosen by an individual and the dependent variable is the price paid. Controlling for plan fixed effects, we can measure how the price paid for a fixed bundle of plans changes across age thresholds. Columns 1 and 2 of Appendix Table A.1 runs these regressions for the July-Dec. 2009 sample and the full sample. Results are similar; the sole exception is the price increase at age 40, which is estimated to be 14-15% in the July to December sample, but is only 9% in the full sample. Columns 3 and 4 estimate equation 2 on both samples and show similar results.

This nonstructural approach shows that there is little response in individual choice to price increases: the increase in spending is approximately equal to the increase in prices. Yet Section 4 will show that these results do *not* imply that individuals are not sensitive to price. Rather, individuals are al-

[20] Linear age splines have knots at each age threshold. Additional controls include indicators for month, geographic region, and gender.

[21] An age group-specific insurance spending elasticity could be estimated by as $\eta_s = \frac{\kappa_s}{\pi_s}$, but the standard errors on each η_s would be extremely large.

ready gravitating to the least generous tier and cheapest plans available. Thus, despite the wide range of plan generosities available in the Connector, individuals do not have much latitude to respond to a price increase. They are simply unable to substitute to cheaper plans. Thus, these results highlight the importance of context in determining the effect of policy changes (e.g., altering the tax exclusion for employer-provided health insurance) and motivate our structural model of consumer preferences in Section 4.

4 Discrete-Choice Model

4.1 Theoretical Model

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

$$u_{ij} = \delta_j + \mu_{ij} + \xi_{ij} + \varepsilon_{ij},$$

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

$$s_{ij} = \frac{\exp\left(\delta_j + \mu_{ij} + \xi_{ij}\right)}{1 + \Sigma_j \exp\left(\delta_j + \mu_{ij} + \xi_{ij}\right)},$$

where s_{ij} represents the probability that consumer i purchases product j. In the absence of individual heterogeneity μ_{ij}, the δ_j parameters simply represent an inversion of the observed market shares for each plan. The mean utility δ_j can be a plan fixed effect, or can be a function of plan characteristics X_j, such as insurer (brand), price, deductibles, and copayments. A given insurance plan (e.g., HMO Blue Basic Value) is offered in multiple markets.

As discussed in the previous section, we identify the premium coefficient 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.[22] More formally, let δ_{j30} be the mean utility of product j offered to consumers who are age 30, where $\delta_{j30} = X_j'\beta + \alpha p_{j30}$, and X is a vector of plan fixed effects. Similarly, let δ_{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:

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

Of course, there are age trends in preferences. We allow for δ_j 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 (with six knots) for each of the 21 different plans would allow maximal 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. 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.

4.2 Estimation

The model is estimated using a conditional logit approach in Table 3. This allows us to interact consumer characteristics, such as age, with plan characteristics, such as price. A critical feature of demand in this market is consumer heterogeneity in preferences. 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

[22] 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 critical in Table 2.

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. Specifically, the unobserved product characteristic, denoted above by ξ_{ij}, consists of age and geographic specific deviations from mean plan quality - the added utility that a specific plan brings to an older consumer, for example. These are likely to be positively correlated with price, and failing to account for such effects will bias the price sensitivity toward zero.[23] 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 3 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 3 cuts the age span in the Connector 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 3 maps out these price elasticities under various specifications, highlighting the pattern in semi-elasticities over the life cycle.

Table 4 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 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

[23] Furthermore, later in this section, we will discuss how to use the identifying variation from coarse insurer pricing.

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, a 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 sensitivity 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.

Taken together, the results from Tables 3 and 4 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, a higher preference for more generous coverage, or both.

Various demographic factors could be driving the preference heterogeneity 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. 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.

5 Application: Age-Based Pricing Regulation

5.1 Motivation

In this section, we model how age-based pricing regulation affects markets in the presence of imperfect competition. Existing work (e.g., Blumberg, Buettgens, and Garrett 2009) has assumed that insurers price differentially by age solely due to cost differentials, so age-based pricing regulations only bind to the extent the ratio of costs exceeds the maximum allowable ratio of prices by age. However, Section 4 shows that price-sensitivity varies by age, and so insurers would want to price discriminate and charge higher prices to older enrollees, even if health costs did not vary by age. This section analyzes the impact of age-based pricing regulations in the presence of 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 the ratio of the highest price to lowest price cannot exceed θ.
- 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 Connector, $\theta = 2$, while PPACA requires states to set $\theta \leq 3$.

5.2 Model and Theoretical Predictions

Consider two types of consumers, old and young, who are purchasing an insurance plan.[24] Costs rise with age, so the old have an average cost (to the

[24] 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. That is, we assume the mandate

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[25] 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.[26] Stricter limitations on age-based pricing (reducing θ) transfers resources from old individuals to young individuals. Age-bands are only binding up to the ratio of costs between the two groups. Prices 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 *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. The existence of the price discrimination motive means that low-cost consumers could face higher prices than high-risk consumers if they are not sensitive to the price of insurance policies

is effective. Section 6 examines noncompliance with the mandate.

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

[26] 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.

(the "worried-well"; see Starc 2012). However, it may also be the case that the higher prices high-risk consumers face are amplified by low price sensitivities.

Now let the market be imperfectly competitive, with all plans being identical in average quality. 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:

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

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. 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 lower the average cost. 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.

5.3 Counterfactual Exercise: Changing Age-Based Pricing Regulation

In this section, we examine how alternative age-based pricing regulations would affect prices and welfare on the Massachusetts Connector. We use the stylized pricing rule developed in Proposition 1, but we do not deny that firm strategies may in fact be much more complicated.[27] However, our framework matches how firms set prices on the Connector in practice. Insurers submit quotes to the Connector, which consist of a base rate and an adjustment factor. This adjustment factor takes into account any differences in costs and may account for differences in consumer preferences.

We consider two types of counterfactual scenarios. In the first one, firms price to two age groups (over/under 45), and in the second, firms price to three age groups (27-35, 35-45, 45+). Before examining age-based differences in costs, we simulate how insurers would price in a stylized environment where a risk-adjustment scheme perfectly compensated insurers for age-based differences in costs. (There is no risk adjustment in the Connector.) When firms are unconstrained by age-based pricing regulation (Age-Unconstrained), young consumers are sufficiently price sensitive to ensure price competition on their policies. Our estimates indicate that the optimal markup for consumers under 45 is 10%. However, because there is a sharp drop in price sensitivity at age 45, the older group of consumers faces high markups (roughly a third of

[27] For instance, we do not capture competition over multiple, related products. For example, a firm may target a specific group of consumers with one plan (i.e., young individuals or families with a bronze plan), and another group with a different plan.

price) in the absence of regulation. Now consider the other extreme: Age-Pooling, so firms can set only one price for all ages. The presence of younger consumers in a pool with older consumers can partially hold down markups, and because they are relatively more prevalent in the population, their higher elasticity has a large weight on the average markup, which is approximately 20%. Thus, moving from Age-Unconstrained pricing to Age-Pooling would entail substantial redistribution away from younger individuals.

These simulations show that older consumers would be willing to pay younger consumers to face premiums set for the entire population, even if risk adjustment perfectly compensated insurers for age-related differences in costs.[28] Table 5 shows the transfers from younger individuals that would result from moving from Age-Unconstrained to Age-Pooling or Age-Bands. When firms price to two age groups, pooling leads to a 6% average increase in premium for those under 45, while if they price to three groups the transfer increases to 7.6% (column 2).[29] As the age-bands are relaxed, the transfers away from the under-45 year olds is correspondingly reduce. Note that these transfer estimates *assume no cost differences* between the two groups; any cost difference would exacerbate this transfer from younger consumers to older consumers.

We next examine differences in costs by age-group. Table 6 uses data from

[28] The desire to face younger consumers' prices arises because they are more price-sensitive and does not rely on younger consumers having lower costs. Note that there could also be different optimal markups over classes of products as well. For example, if consumers purchasing bronze plans are more price-sensitive than average, the average markup on these plans will be lower. In addition, the minimum price effect (if it is truly a heuristic) will induce plans to compete vigorously to be the cheapest. None of these effects are explicitly modeled here, however additional specifications (not presented) indicate that the price sensitivity for bronze plans is significantly higher than their silver or gold counterparts.

Note that a dollar reduction in premium to a younger consumer is more valuable in utility terms than the corresponding dollar increase in premium to an older consumer. The only potential welfare gains come from more efficient sorting of consumers. We abstract from consumer reoptimization.

[29] The simulation results are somewhat sensitive to a number of choices, including which variation to use (for example, identifying the price sensitivity of a 31-year-old using the price jump at 30 or the price jump at 35), the flexibility of the demand specification, and the weighting placed on different age groups. Therefore, the results in each simulation are always similar in direction and order of magnitude, though not identical.

the 2008 Medical Expenditure Panel Survey (MEPS) on the health costs of difference 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.[30]

We find that 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 modified community rating $\theta = 2$ 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 discrimination 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.

Increasing the age-band θ reduces the transfer that young consumers give to older consumers. The level of θ is subject to regulation, and varies: while Massachusetts has imposed a 2:1 age band, PPACA calls for a 3:1 age band. Table 5 additionally shows how transfers change with modified community rating rules. The PPACA regulation lowers prices for consumers under 40 by approximately 3% relative to the Massachusetts regulation, before any cost differences are taken into account. Another way of describing the impact of preference heterogeneity is to consider its impact on regulation.

Any benefit of providing younger consumers with a separate market with

[30] 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.

different types of plans should be weighed against the premium-reducing impact these consumers would have in the broader market, both from lower medical expenditures and higher price sensitivity. This is relevant not only when considering pooling across ages, but pooling across incomes as well. In Massachusetts, the subsidized market is separated from the unsubsidized market, but this need not be the case. If the markets were combined, assuming the subsidized consumers are more price sensitive, premiums would be lower for all of the unsubsidized consumers, leading to a redistribution from the subsidized consumers to unsubsidized consumers.

The discussion of age-based pricing regulation assumes that age-bands will bind if θ is less than the ratio of costs in the data. However, in the presence of imperfect competition and preference heterogeneity, regulation will bind for ratios of costs that are much lower than θ. We note that even smaller cost differences lead to binding regulation: for example, if insurers coarsely price to just 3 age groups, a θ of 2 will bind even if the ratio of costs is only 1.5, and θ of 3 will bind when the ratio of costs is only 2.5. This indicates that, given the cost differences in the MEPS, the Massachusetts age bands of 2:1 will certainly bind, and the federal age bands of 3:1 are likely to bind.

Finally, Figure 4 plots premiums by age assuming the costs in the MEPS data, the preferences in the demand system, and no regulation, under both perfect and imperfect competition. First, we note that there is a dramatic increase in costs at age 55. However, differences in preferences amplify the differences in costs, leading to a much larger gap between prices under perfect and imperfect competition. Furthermore, we not that the prices under imperfect competition are more extreme than those we observe in the data, which is reasonable given that they are unconstrained by regulation in this simulation. Finally, 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.

6 Market Participation and Markups

Our results are likely to represent a lower bound on the potential impact of modified community rating on younger consumers, as we have assumed all consumers purchase insurance–i.e. that the mandate is effective. In the absence of a mandate, consumers may opt out of coverage. If in the face of higher prices, younger consumers simply opt out of the market completely, this will exacerbate the transfers from the younger consumers left in the market to older consumers. The model above can be expanded to allow consumers to opt out of the market; for simplicity, we consider the full pooling case. Denote the participation rate of the young consumers by c_L and their take-up elasticity by ε_L. Then μ_L represents the percent markup under unconstrained pricing and $\overline{\mu}$ the percent markup under full pooling. The participation rate can then be written as

$$c_L = 1 - \varepsilon_L \left(\mu_L - \overline{\mu}\right).$$

The participation rate can be defined similarly for the older consumers. The optimal pooled markup can then be expressed as

$$\overline{\mu} = \frac{\sigma c_L \mu_L + (1-\sigma)\, c_H \mu_H}{\sigma c_L + (1-\sigma)\, c_h}.$$

We simulate the effect of allowing for opting out in the absence of a mandate in Figure 5. Simulation 1 (solid line), assumes full market participation among the older group. Using intermediate values from the previous simulations, we use a markup of 15% for the younger consumers and a percent markup of 35% for the older 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, as noted earlier in the paper, 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 compliance) to -5. We use optimal markups for each group as estimated above, as well as the empirical population shares. The relationship is nonlinear, but for a take-up elasticity of -5, no younger consumers participate in the market and the optimal pooled markup is equal to the optimal markup for older consumers. Therefore, we can get a so-called death spiral from differences in preferences alone. Furthermore, to the extent that elasticities are correlated with underlying costs, imperfect competition can amplify the potential for a death spiral. In either case, our estimates of transfers are certainly a lower bound given the potential for less than full market participation.

Our results are more general and apply to any situation in which ε_L is greater in absolute magnitude than ε_H. Simulation 2 (dashed lines) represent a simulation in which both older and younger consumers are allowed to have nonzero take-up elasticities; however, the elasticity of the younger consumers is constrained to be exactly twice that of the younger consumers. So long as the older consumers are less responsive to price than the younger consumers, the death spiral pattern holds.[31] 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 using a novel data set in a health insurance exchange that servers as a model for national health reform. Our analysis of choices speaks to a number of policy and conceptual issues. First, calculating the insurance spending elasticity depends critically on the choice set available to consumers. Therefore, we use a discrete choice approach to

[31] To the extent that elasticities are correlated with underlying costs, the effects of selection are likely to be similar. In addition, the general effect of nonzero take-up elasticities for older consumers is to increase the numbers of older consumers in the market, as full pooling represents a price decrease for older consumers.

measure consumers' price sensitivity, and find broad evidence of consumer heterogeneity in preferences. A consumer in the 75th percentile of the price sensitivity distribution is four times more price sensitive than a counterpart in the 25th percentile.

Modified community rating and age-based pricing in the Commonwealth Choice program are crucially important regulations. Priceable variation in consumer preferences gives insurers an additional motive, due to price discrimination, to increase premiums to older consumers. Consistent with findings in the insurance literature, we find that consumers have very heterogeneous preferences. We extend these findings to argue that imperfect competition may amplify variation in insurance prices if inelastic consumers also tend to be high cost consumers. Therefore, regulators should be especially cautious in defining risk pools, as heterogeneous preferences for insurance can exacerbate differences in premiums among different groups of consumers.

Beyond identifying crucial regulatory features in HIEs, we also demonstrate *why* these regulations are important. Exchange designers must not only consider the nature of consumer demand, but also strategic insurer pricing in the face of both consumer demand and the regulatory regime. Understanding behavior on the exchanges is crucial to designing them well. Our results from Massachusetts provide an early look at HIE, and a foundation for researchers who study health exchanges and policymakers who design such exchanges.

References

Abaluck, J. and J. Gruber. 2011. Heterogeneity in Choice Inconsistencies among the Elderly: Evidence from Prescription Drug Plan Choice. The American Economic Review 101 (3): 377–381.

Blumberg, L.J., M. Buettgens, and B. Garrett. 2009. Age Rating Under Comprehensive Health Care Reform: Implications for Coverage, Costs, and Household Financial Burdens. Urban Institute Timely Analysis of Immediate Health Policy Issues. Available at: http://www. urban. org/UploadedPDF/411970_age _rating. pdf.

Buchmueller, T.C. and J. DiNardo. 1999. Did Community Rating Induce an Adverse Selection Death Spiral? Evidence from New York, Pennsylvania and Connecticut.

Bundorf, K., J. Levin, and N. Mahoney. 2008. Pricing, Matching and Efficiency in Health Plan Choice. mimeograph, Stanford University.

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

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

Cutler, D.M. and S.J. Reber. 1998. Paying for Health Insurance: The Trade-Off between Competition and Adverse Selection. Quarterly Journal of Economics 113 (2): 433–466.

Dafny, L., M. Duggan, and S. Ramanarayanan. 2009. Paying a premium on your premium? Consolidation in the US health insurance industry. National Bureau of Economic Research.

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

Duggan, M., P. Healy, and F.S. Morton. 2008. Providing prescription drug coverage to the elderly: America's experiment with Medicare Part D. The Journal of Economic Perspectives 22 (4): 69–92.

Einav, L., A. Finkelstein, and M.R. Cullen. 2010. Estimating Welfare in Insurance Markets Using Variation in Prices. Quarterly Journal of Economics 125 (3): 877–921.

Ericson, K.M.M. 2010. Market Design when Firms Interact with Inertial Consumers: Evidence from Medicare Part D.

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

Geruso, M. 2011. Community Rating in Employer Health Insurance: Inefficiencies Beyond Adverse Selection.

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

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

Herring, B. and M.V. Pauly. 2006. Incentive-compatible guaranteed renewable health insurance premiums. Journal of Health Economics 25 (3): 395–417.

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.

Kaufman, D.W., J.P. Kelly, L. Rosenberg, T.E. Anderson, and A.A. Mitchell. 2002. Recent patterns of medication use in the ambulatory adult population of the United States. JAMA: The Journal of the American Medical Association 287 (3): 337.

Ketcham, J.D., C. Lucarelli, E.J. Miravete, and M.C. Roebuck. 2010. Sinking, swimming, or learning to swim in Medicare Part D. American Economic Review.

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

Krueger, A.B. and I. Kuziemko. 2011. The Demand for Health Insurance among Uninsured Americans: Results of a Survey Experiment and Implications for Policy.

Long, S. and L. Phadera. 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.

Lustig, J. 2010. The welfare effects of adverse selection in privatized Medicare. Department of Economics, Boston University.

Roe, C.M., A.M. McNamara, and B.R. Motheral. 2002. Gender-and age-related prescription drug use patterns. The Annals of pharmacotherapy 36 (1): 30.

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

Starc, A. 2010. Insurer Pricing and Consumer Welfare: Evidence from Medigap.

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

Zuckerman, S. and S. Rajan. 1999. An alternative approach to measuring the effects of insurance market reforms. Inquiry: a journal of medical care organization, provision and financing 36 (1): 44.

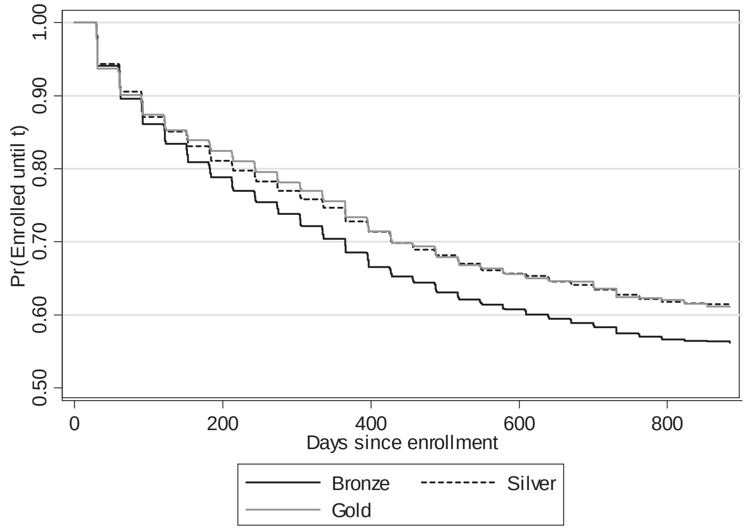

Figure 1: Length of Time Enrolled in the Connector by Tier of Plan. Notes: Kaplan-Meier survival estimate. Sample: all individuals, July 2007 to Dec 2009.

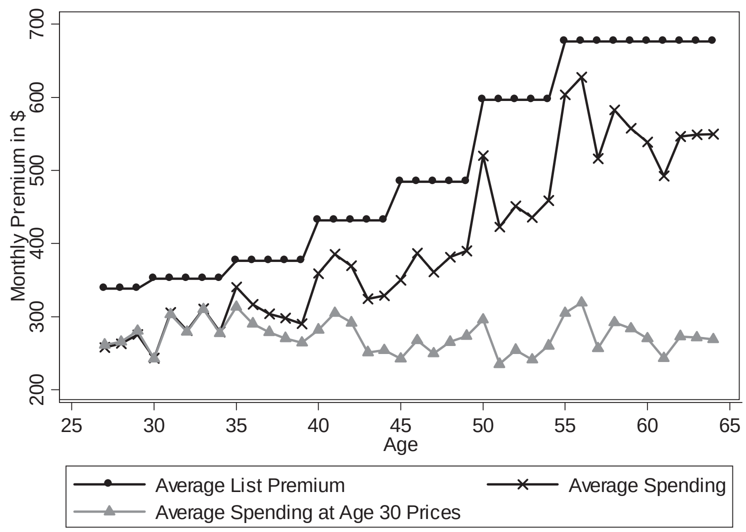

Figure 2: Average Prices and Spending, By Age. Notes: Data: Nov. and Dec. 2009. Average list premium is the plan-weighted average using the Nov. plan-zipcode price. Average spending is person-weighted. Average spending at age 30 prices uses actual choices but prices for a 30-year old in Nov. in that zipcode.

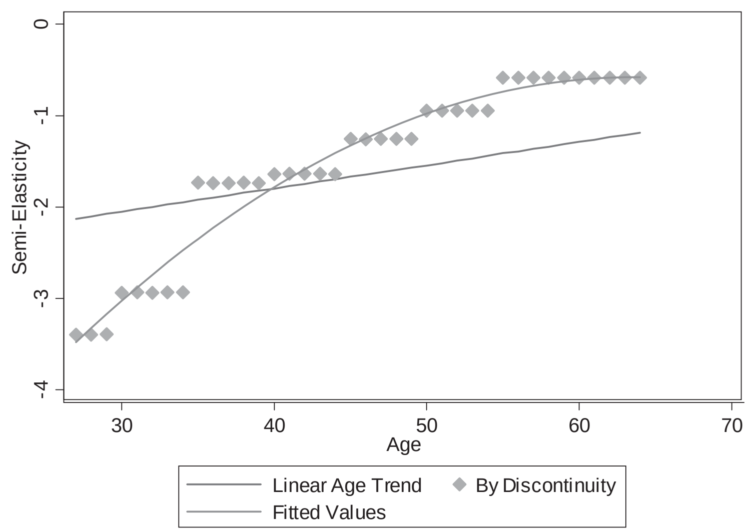

Figure 3: Semi-Elasticities By Age. Notes: Semi-elasticity describes behavioral response (reduction in market share) to a $100 increase in monthly premium. Linear age trend is plotted using results from Column 3 of Panel A of Table 3. Discontinuity results plot the average semi-elasticity obtained from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

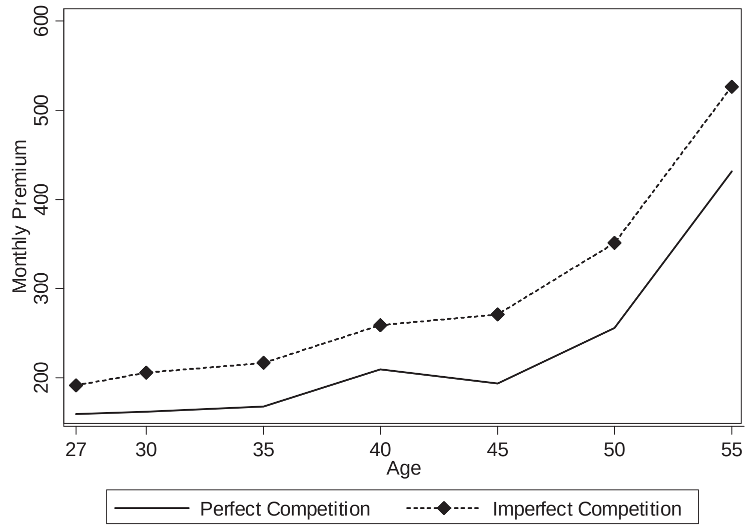

Figure 4: Simulated Monthly Premiums With Unconstrained Pricing. Notes: Assumes no age-based pricing regulation. Under perfect competition, we assume insurers charge at cost, with costs taken from the 2008 MEPS (see Table 6). Under imperfect competition, we assume firms charge the optimal markup based on age-specific price sensitivity calculated in Table 3 Panel B.

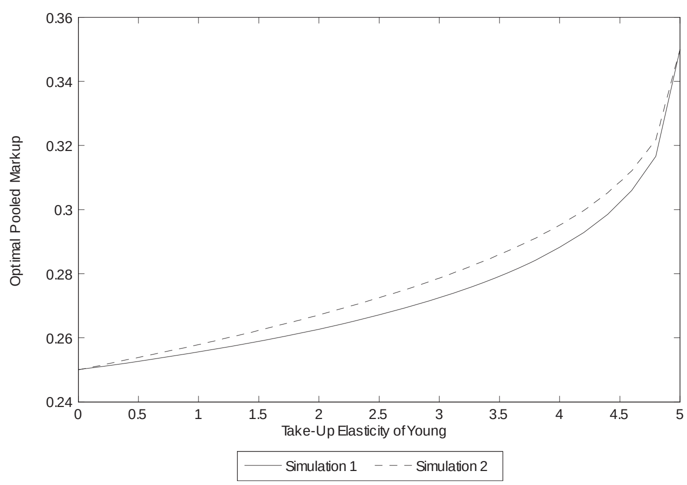

Figure 5: Optimal Markups in the Absence of a Mandate. Notes: Based on Simulation 1 and 2, described in the text.

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 and Spending Response to Age Discontinuities

	ln(price index)	ln(premiums paid)
Indicators:		
Above 30	0.0224***	-0.0438
	(0.00172)	(0.0323)
Above 35	0.0790***	0.0442
	(0.00199)	(0.0411)
Above 40	0.150***	0.147***
	(0.00217)	(0.0447)
Above 45	0.106***	0.0138
	(0.00189)	(0.0440)
Above 50	0.204***	0.207***
	(0.00201)	(0.0502)
Above 55	0.128***	0.192***
	(0.00232)	(0.0462)
Linear Age Spline	Yes	Yes
Basic Controls	Yes	Yes
N Persons	2,616	2,616
R^2	0.998	0.572
	IV-Stage 1 from Column 1	
		ln(premiums paid)
	ln(price index)	0.962
		(0.176)
	Linear Age Spline	Yes
	Basic Controls	Yes
	R^2	0.569

Sample: July-Dec 2009. Note: Heteroskedasticity robust standard errors in parentheses. Age spline consists of piecewise linear age controls within each age group. Controls include indicators for month of enrollment, indicators for geographic market, and gender. IV results from two-stage least squares. *** p<0.01, ** p<0.05, * p<0.1.

Table 3: 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 4: 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	Carrier	Carrier
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 5: Transfers From Under 45 Year Olds, as a Percent of Premiums

	Firms Price to:	
	2 Age Groups	3 Age Groups
Full Pooling	6.01%	7.55%
2:1 Age Bands	4.80%	5.98%
3:1 Age Bands	1.43%	3.20%

Notes: Transfers are calculated as the differences in optimal firm markups under each regulatory policy. Transfers are calculated assuming constant costs across ages.

Table 6: 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 in the New England states. Avg. insurer cost is mean priviate insurer expenditure for this sample.