Equilibria in Health Exchanges: Adverse Selection vs. Reclassification Risk

Ben Handel[*], Igal Hendel[†], and Michael D. Whinston[‡§]

August 21, 2013

Abstract

This paper studies regulated health insurance markets known as exchanges, motivated by their inclusion in the Affordable Care Act (ACA). We use detailed health plan choice and utilization data to model individual-level projected health risk and risk preferences. We combine the estimated joint distribution of risk and risk preferences with a model of competitive insurance markets to predict outcomes under different regulations that govern insurers' ability to use health status information in pricing. We investigate the welfare implications of these regulations with an emphasis on two potential sources of inefficiency: (i) adverse selection and (ii) premium reclassification risk.

We find that market unravelling from adverse selection is substantial under the proposed pricing rules in the Affordable Care Act (ACA), implying limited coverage for individuals beyond the lowest coverage (Bronze) health plan permitted. Although adverse selection can be attenuated by allowing (partial) pricing of health status, our estimated risk preferences imply that this would create a welfare loss from reclassification risk that is substantially larger than the gains from increasing within-year coverage, provided that consumers can borrow when young to smooth consumption or that age-based pricing is allowed. We extend the analysis to investigate some related issues, including (i) age-based pricing regulation (ii) exchange participation if the individual mandate is unenforceable and (iii) insurer risk-adjustment transfers.

[*]Department of Economics, UC Berkeley; handel@berkeley.edu

[†]Department of Economics, Northwestern University; igal@northwestern.edu

[‡]Sloan School and Department of Economics, M.I.T.; whinston@mit.edu

[§]We thank conference discussants Gautam Gowrisankaran, Bruno Jullien, Kei Kawai, Pierre-Thomas Leger, and Neale Mahoney for their detailed advice and comments on this paper. We also thank the participants in seminars at AEA Annual Meetings (2012), Berkeley, Berkeley-Stanford IO Fest (2011), Bureau of Economic Analysis, Carnegie Mellon Heinz, Duke Applied Microeconomics Jamboree (2012), Harvard, Montreal HEC, Montreal HEC Health-IO Conference, M.I.T., NBER Health Care Summer Institute (2013), NYU, Northwestern-Toulouse IO Conference, Robert Wood Johnson SHPR Annual Meeting (2012), Stanford SITE: Theory-Based Modeling (2012), Toulouse Network for IT Annual Meeting (2011), University of Arizona, University of Chicago, UCL, University of Wisconsin-Milwaukee, Utah Winter Business Economics Conference (2013), Wharton Health Exchanges Conference, and Yale. All authors are grateful for support from NSF grant SES-1259770 and Whinston also thanks prior support from NSF and the Toulouse Network for Information Technology. We thank Jorge Lemus and Fernando Luco for outstanding research assistance.

1 Introduction

Health insurance markets almost everywhere are subject to a variety of regulations designed to encourage the efficient provision of insurance. In the United States, the Affordable Care Act (ACA), passed in 2010, defines a class of regulated state-by-state markets, called *exchanges,* in which insurers must offer annual policies that comply with specific federal rules related to insurance contract design and pricing. Relative to the private individual insurance markets that existed prior to the Affordable Care Act, the exchanges remove almost entirely the ability of insurers to price based on consumers' health conditions, heavily regulate both financial and non-financial contract dimensions, require all individuals to have insurance, and organize and present information with the intent to facilitate informed consumer decision-making.[1] While state exchange regulators will have some flexibility to implement their own market designs, all exchanges will share these common characteristics when they begin, or continue to, insure consumers in 2014. While there has certainly been a great deal of public discussion concerning the desirability of this reform, there has been surprisingly little formal analysis of the likely outcomes in these exchanges and the welfare impacts of alternative designs that were considered but not implemented.

This paper sets up and empirically investigates a model of insurer competition in a regulated marketplace, motivated by these exchanges. We focus on the issue of premium regulation and ask how different insurer pricing restrictions would impact consumer welfare. Specifically, we start with pure community rating as a default, and then investigate a range of alternative regulations that allow insurers greater flexibility in pricing individual-specific characteristics such as pre-existing medical conditions.[2] Relative to these alternative regulations, the ACA prohibition on pricing nearly all pre-existing conditions can directly impact two distinct determinants of consumer welfare: adverse selection and re-classification risk.[3] Adverse selection is present when there is individual-specific information that can't be priced, and sicker individuals tend to select greater coverage.[4] Reclassification risk, on the other hand, arises when insurance contracts are of limited duration and changes in health status lead to changes in premiums. In our setting, reductions in the extent to which premiums can be based on pre-existing conditions are likely to increase the extent of adverse selection, but reduce the reclassification risk that insured individuals face. For example, when pricing based on pre-existing conditions is completely prohibited (which is close to the case in the current regulation), reclassification

[1] For example, in all states, insurers must offer the same premium to different indivduals of the same age (subject to some minor caveats), and premiums to individuals of different ages cannot differ by more than a 3:1 ratio. Federal regulations govern the minimum actuarial standards for contracts nationwide, while states have some leeway both to further restrict these financial standards and to determine what medical procedures insurers must cover. As we discuss below, the ACA also bans pricing based on nearly all pre-existing conditions.

[2] See, e.g., Bhattacharya et al. (2013) or Capretta and Miller (2010) for policy-oriented discussions that advocate relaxing the pricing restrictions present in the ACA (subject to some complementary market design changes).

[3] Each of these phenomena is often cited as a key reason why market regulation is so prevalent in this sector in the first place.

[4] See Akerlof (1970) and Rothschild and Stiglitz (1976) for seminal theoretical work.

risk is eliminated but adverse selection is likely to be present. At the other extreme, were unrestricted pricing based on health status allowed, adverse selection would be completely eliminated. We would then expect efficient insurance provision conditional on the set of available contracts, although at a very high price for sick consumers.[5] Thus, in determining the degree to which pricing of pre-existing conditions should be allowed, a regulator needs to consider the potential trade-off between adverse selection and re-classification risk.

To study the impact of the ACA and alternative regulations we develop a stylized model of an insurance exchange that builds on work by Rothschild and Stiglitz (1976), Wilson (1977), Miyazaki (1977), Riley (1985) and Engers and Fernandez (1987) who all modeled competitive markets with asymmetric information. In the model, the population is characterized by a joint distribution of risk preferences and health risk and there is free entry of insurers. We assume that individuals are forced to buy insurance in the marketplace, as a result of a fully-enforced individual mandate (we relax this in an extension). Throughout the analysis, we fix two classes of insurance contracts that each insurer can offer.[6] The more comprehensive contract has 90% actuarial value and mimics the most generous coverage allowed under the ACA, while the less comprehensive contract has 60% actuarial value and mimics the least generous coverage allowed under the ACA.[7] These contracts are required to be annual, as in the current legislation. The model abstracts away from horizontal insurer differentiation from, e.g., different access to medical providers and treatments.[8]

The challenges in conducting this analysis are both theoretical and empirical. From the theoretical perspective, the analysis of competitive markets under asymmetric information, specifically insurance markets, is delicate. Equilibria are difficult to characterize, can be sensitive to the contracting assumptions, and are often fraught with non-existence. On the empirical side, any prediction of exchange outcomes must naturally depend on the extent of information asymmetries, that is, on the distribution of risks and the information in the hands of insurees. Thus, a key empirical challenge is identifying these distributions.

To deal with the Nash equilibrium existence problems highlighted by Rothschild and Stiglitz (1976) we focus on another concept developed in the theoretical literature: Riley equilibria [Riley (1979)]. Under the Riley notion, firms consider competitors' reactions so that deviations rendered unprofitable by subsequent reactions are not undertaken. The main roles of the theoretical analysis are (i) to prove the existence and uniqueness of Riley equilibrium in our context and (ii) develop algorithms to find

[5] This abstracts away from liquidity concerns that could be present in reality, especially for low income populations.

[6] Importantly, our model allows insurers to offer both kinds of insurance contracts simultaneously. In the ACA, insurers are required to offer at least two policies, in the 80% (gold) and 70% (silver) actuarial equivalence classes [see, e.g., Fernandez and Mach (2012)].

[7] Actuarial value reflects the proportion of total expenses that an insurance contract would cover if the entire population were enrolled. In addition to the contracts we study, the ACA permits insurers to offer two classes of intermediate contracts with 70% and 80% actuarial value respectively. In the legislation, 90% is referred to as "platinum", 80% "gold", 70% "silver", and 60% "bronze."

[8] While such horizontal differentiation could be important for choice / pricing in practice, here we focus on the financial role of insurance in risk protection and the subsequent trade-off between adverse selection and reclassification risk.

both the Riley equilibrium and any Nash equilibria, should they exist.[9] As an extension in Section 6, we study an alternative equilibrium notion, Wilson equilibrium [Wilson (1977)], which places different restrictions on possible equilibrium deviations.

We use the outputs of this equilibrium market analysis (premiums and consumers' plan choices) as inputs into a long-run welfare model that integrates year-to-year premium risk, conditional on the pricing regulation and underlying health transition process. This model evaluates welfare from the perspective of an ex-ante unborn individual, and follows an individual through many consecutive one-year markets characterized by the static model. We evaluate lifetime welfare under two different scenarios. On the one hand, we consider fixed income over time, which is a reasonable assumption when borrowing is feasible. Alternatively, to capture potential borrowing frictions, we also evaluate welfare under the observed income profile. One benefit of pricing health conditions is that the population is healthier at younger ages, when their income is lower. Pricing pre-existing conditions, which results in lower premiums early in life, can therefore be beneficial for steep enough income profiles.

To simulate the market we need a population of potential insureds for whom we know the joint distribution of ex ante health status and risk preferences. We obtain this information using individual-level health plan choice and health claims data for the employees of a large firm and their dependents. We leverage several unique features of the data to cleanly identify risk preferences including (i) a year where all employees made active, non-default choices, due to a menu change and (ii) the fact that the plans available differ financially, but not in terms of provider availability. We develop a structural choice model, that generalizes Handel (2013), to quantify risk preferences.[10] In particular, we estimate a distribution of heterogeneous risk preferences that is allowed to depend on an individual's ex ante health status since prior work on insurance markets [see e.g. Finkelstein and McGarry (2006) or Cohen and Einav (2007)] reveals that correlation between health risk and risk preferences can have important implications for market outcomes (e.g., the extent of adverse selection).

To model health risk perceived by employees at the time of plan choice, we use the methodology developed in Handel (2013), which characterizes both total cost health risk and plan-specific out-of-pocket expenditure risk. The model incorporates past diagnostic and cost information into individual-level and plan-specific expense projections using both (i) sophisticated predictive software developed at Johns Hopkins Medical School and (ii) a detailed model of how different types of medical claims translate into out-of-pocket expenditures in each plan. This cost model outputs an individual-plan-time specific distribution of predicted out-of-pocket expenditures that we incorporate into the choice model under the assumption that consumer beliefs about future health expenditures conform to our cost model estimates.

We use the estimates to study market equilibria and long-run welfare in counterfactual market

[9] We study Nash equilbria for two cases: (i) when insurers can offer both policies at once and (ii) when insurers are restricted to offering only one policy.

[10] While we incorporate consumer inertia in estimation to correctly estimate risk preferences, as in Handel (2013), our subsequent exchange equilibrium analysis studies a static marketplace where consumers make active non-inertial choices.

environments. While we realize that our sample, coming from one large firm, is not an externally valid sample on which to base a policy conclusion, the depth and scale of the data present an excellent opportunity to illustrate our methodology. In Section 6, we re-run our main analyses with a re-weighted sample that matches our population to that in the nationally representative Medical Expenditure Panel Survey (MEPS) on key demographic dimensions. This analysis yields, in general, similar results to those from our primary sample.

For the static market with pure community rating (no price discrimination) our results show substantial within-market adverse selection. The Riley equilibrium results in full unravelling, with all consumers purchasing a 60% plan at a premium equal to plan average cost for the entire population. There is still full unravelling in each age cohort once we allow for age-based pricing, though the premiums for each age group reflect the plan average costs conditional on age. High-risk consumers have large price externalities on healthier ones, leading to extremely expensive, and essentially unavailable, 90% plans. This suggests that the Minimum Creditable Coverage in terms of actuarial value, regulated to 60% in the ACA, could be a pivotal determinant of consumer welfare in the exchanges.[11]

We study alternative policies where insurers can price pre-existing conditions to some extent. For illustrative purposes, we consider the case where insurers can price based on ex ante health status quartiles: here the Riley equilibria across quartiles result in less adverse selection in the sense that both the 60% and 90% plans have positive market share (though for some quartiles the market still fully unravels). We then study pricing based on finer partitions of health-status, all the way up to the case of full risk-rating, where insurers can use all available information to price policies. As insurers can price on more and more health-relevant information the market share of consumers enrolled in the 90% policy increases, implying reduced adverse selection. In all cases, Nash equilibria coincide with the Riley outcomes if firms offer only one policy [as in Rothschild and Stiglitz (1976)], while for most cases Nash equilibria fail to exist when firms can offer multiple policies.

Though greater ability to price health-status information reduces adverse selection, our long-run welfare results illustrate the extent to which such policies exacerbate reclassification risk. Under the case of fixed income from age 25 to 65, welfare is highest when health-status pricing is banned. For example, from an ex ante perspective an individual with median risk aversion would be willing to pay $3,082 each year from 25-65 to be in a market with pure community rating relative to the case of pricing based on health-status quartiles, though the latter yields greater within-year coverage. This is approximately five times the $619 welfare loss that occurs from adverse selection under pure community rating (conditional on the restricted set of available contracts), and roughly half of the average annual medical expenses in the population. Thus, the welfare losses due to reclassification risk, even for fairly limited pricing of health status, can be quantitatively large. Moreover, as the ability to price on health-status becomes greater, the welfare loss becomes larger. Finally, when we change the fixed

[11] Interestingly, the market unravelling we find under community rating (with or without age-based pricing) is somewhat consistent with experience in the Massachusetts exchange, where most buyers opted for the Bronze (60%) plan in the early years of this ACA-like exchange [see, e.g., Ericson and Starc (2013)].

lifetime income assumption and allow for increasing income profiles the losses from reclassification risk are attenuated because health-status based pricing decreases premiums earlier in life when income is lower (and thus serves the purpose of smoothing income over time). This effect is eliminated, however, if age-based pricing is allowed (as in the ACA).

We study several extensions to address issues of particular interest under the ACA. In addition to investigating age-based pricing, discussed above, we use our framework and estimates to quantify the subsidies necessary to guarantee different levels of participation in the exchange. This links directly to the issue of whether individuals will adversely select into the exchange based on health status. We find that, absent subsidies or penalties, approximately 26% of the population would opt out of the exchange when the pricing of pre-existing conditions is prohibited. Those who opt out are mostly younger and healthier individuals: about half of the 30 to 35 year old population would prefer to opt out. As the healthier types opt out premiums increase, leading to further desertions. With no subsidies or penalties, premiums in the market are approximately 30% higher than in the case of full participation. A subsidy of over $3,000 per person/year is required to decrease the percentage opting out to 10%.[12]

We also study an extension that allows for risk-adjustment transfers between insurers, as stipulated in the ACA. These transfers are designed to subsidize insurers who take on higher risks and, consequently, ameliorate adverse selection. We use the model to evaluate the impact of the adjustment formula proposed by the Federal government [see, e.g., Dept. of Health and Human Services (2012a) or Dept. of Health and Human Services (2012b)]. While in practice risk adjustment can lead to a number of problems, such as insurers up-coding enrollees to qualify for larger transfers, we abstract from such issues and assume that the government can perfectly observe the health status of each enrollee. The Riley equilibrium with this insurer risk-adjustment has 49% of the population in the 90% policy, as a result of reduced adverse selection.

This paper builds on related work that studies the welfare consequences of adverse selection in insurance markets by adding in a long-term dimension whereby price regulation induces a potential trade-off with re-classification risk. Relevant work that focuses primarily on adverse selection includes Cutler and Reber (1998), Cardon and Hendel (2001), Carlin and Town (2009), Lustig (2010), Einav et al. (2010c), and Bundorf et al. (2012). Handel (2013) and Einav et al. (2013) study the welfare consequences of adverse selection in the contexts of inertia and moral hazard respectively. Ericson and Starc (2013) and Kolstad and Kowalski (2012) study plan selection and regulation in the Massachusetts Connector health insurance exchange. These papers all focus on welfare in the context of a short-run marketplace. There is more limited work studying reclassification risk and long-run welfare in insurance markets. Cochrane (1995) studies dynamic insurance from a purely theoretical perspective in an environment where fully contingent long-run contracts are possible. Herring and Pauly (2006) studies guaranteed renewable premiums and the extent to which they effectively protect consumers

[12] Age-based pricing increases voluntary participation, as younger individuals do not have to subsidize older ones, but on average participation does not increase by much. Only 77% of the population would voluntary participate with age-based pricing without a mandate.

from reclassification risk. Hendel and Lizzeri (2003) and Finkelstein at al. (2005) study dynamic insurance contracts with one-sided commitment, while Koch (2010) studies pricing regulations based on age from an efficiency perspective. Bundorf et al. (2012), while focusing on a static marketplace, also analyze reclassification risk in an employer setting using a two-year time horizon and subsidy and pricing regulations relevant to their large employer context. To our knowledge, there is no similar work that empirically studies the long-run welfare consequences of reclassification risk and adverse selection in an equilibrium setting as a function of price regulation.

The rest of the paper proceeds as follows: In Section 2 we present our model of insurance exchanges and characterize Riley and Nash equilibria in the context of our model. Section 3 describes our data and estimation. In Section 4 we analyze exchange equilibria for a range of regulations on health status-based pricing. Section 5 analyzes the long-run welfare properties of these equilibria, while Section 6 discusses extensions of our main analysis. Section 7 concludes.

2 Model of Health Exchanges

In this section, we describe our health exchange model and provide a set of characterization results. These results provide the algorithm for identifying equilibria using our data, which we do in Section 4.

Throughout the paper, we focus on a model of health exchanges in which two prescribed policies are traded. In our basic specification, these policies will cover roughly 90% and 60% respectively of an insured individual's costs. As such, we refer to these as the "90 policy" and the "60 policy." Within each exchange, the policies offered by different companies are regarded as perfectly homogeneous by consumers; only their premiums may differ. There is a set of consumers, who differ in their likelihood of needing medical procedures and in their preferences (e.g., their risk aversion). We denote by $\theta \in [\underline{\theta}, \overline{\theta}] \subseteq \mathbb{R}_+$ a consumer's "type," which we take to be the price difference at which he is indifferent between the 90 policy and the 60 policy. That is, if P_{90} and P_{60} are the premiums (prices) of the two policies, then a consumer whose θ is below $P_{90} - P_{60}$ prefers the 60 policy, a consumer with θ above $P_{90} - P_{60}$ prefers the 90 policy, and one with $\theta = P_{90} - P_{60}$ is indifferent.

Note that consumers with a given θ may have different underlying medical risks and/or preferences, but will make identical choices between policies for any prices. Hence, there is no reason to distinguish among them in the model. Keep in mind, as we define below the costs of insuring type θ buyers, that those costs represent the expected costs of insuring all of the — possibly heterogeneous — individuals characterized by a specific θ.

Throughout our main specification, we assume that there is an individual mandate that requires that individuals purchase one of the two policies. (But see Section 6.2 for an analysis of participation.)

The costs of insuring an individual of type θ under policy k are $C_k(\theta)$ for $k = 90, 60$. Recall that if the price difference is $\Delta P = P_{90} - P_{60}$, those consumers with $\theta < \Delta P$ prefer policy 60, while those with $\theta > \Delta P$ prefer policy 90. Given this fact, we can define the average costs of serving the populations who choose each policy for a given ΔP to be

$$AC_{90}(\Delta P) \equiv E[C_{90}(\theta)|\theta \geq \Delta P]$$

and

$$AC_{60}(\Delta P) \equiv E[C_{60}(\theta)|\theta \leq \Delta P].$$

We also define the difference in average costs between the two policies, conditional on a price difference $\Delta P \in [\underline{\theta}, \overline{\theta}]$, to be $\Delta AC(\Delta P) \equiv AC_{90}(\Delta P) - AC_{60}(\Delta P)$.

We make the following two assumptions:

Assumption 1: $C_{90}(\theta)$ and $C_{60}(\theta)$ are continuous increasing functions, with $C_{90}(\theta) > C_{60}(\theta)$ for all θ.

Assumption 2: θ has a continuous distribution function F.

The assumption that $C_{90}(\theta) > C_{60}(\theta)$ for all θ simply says that the 90 policy covers more of a consumer's expenses (in expectation) than does the 60 policy.[13] The first part of Assumption 1, on the other hand, is an adverse selection assumption: those consumers who are willing to pay more for the greater coverage in the 90 policy are also the most costly to insure. Since the consumers who choose the 90 policy are those in the set $\{\theta : \theta \geq \Delta P\}$, the assumption implies that

$$AC_{90}(\Delta P) > C_{90}(\Delta P) > C_{60}(\Delta P) > AC_{60}(\Delta P)$$

at any ΔP at which both policies are chosen; i.e., at any $\Delta P \in (\underline{\theta}, \overline{\theta})$, and that $AC_k(\Delta P)$ is increasing in ΔP for $k = 60, 90$. It will also be convenient to define for each policy $k = 60, 90$ the largest and smallest possible average costs: $\underline{AC}_k \equiv AC_k(\underline{\theta})$ and $\overline{AC}_k \equiv AC_k(\overline{\theta})$. Assumption 2 ensures that the function $AC_k(\cdot)$ is continuous for $k = 90, 60$.

In summary, with these assumptions we have the following *Adverse Selection Property* upon which our results will hinge:

Adverse Selection Property $AC_{90}(\theta)$ and $AC_{60}(\theta)$ are continuous monotone functions that are strictly increasing at all $\Delta P \in (\underline{\theta}, \overline{\theta})$, with $AC_{90}(\theta) > AC_{60}(\theta)$ for all θ.

We refer to the lowest prices offered for the 90 and 60 policies as a *price configuration*. We next define the profits earned by the firms offering those prices. Specifically, for any price configuration (P_{90}, P_{60}) define

$$\Pi_{90}(P_{90}, P_{60}) \equiv \left\{ \begin{array}{c} [P_{90} - AC_{90}(\Delta P)][1 - F(\Delta P)] \text{ if } \Delta P \leq \overline{\theta} \\ 0 \text{ if } \Delta P > \overline{\theta} \end{array} \right\}$$

[13] In our empirical work, the 90 policy will in fact dominate the 60 policy in its coverage levels.

and

$$\Pi_{60}(P_{90}, P_{60}) \equiv \left\{ \begin{array}{c} [P_{60} - AC_{60}(\Delta P)]F(\Delta P) \text{ if } \Delta P \geq \underline{\theta} \\ 0 \text{ if } \Delta P < \underline{\theta} \end{array} \right\}.$$

as the aggregate profit from consumers who choose each of the two policies. Let

$$\Pi(P_{90}, P_{60}) \equiv \Pi_{90}(P_{90}, P_{60}) + \Pi_{60}(P_{90}, P_{60})$$

be aggregate profit from the entire population. The set of *break-even price configurations,* which lead each policy to earn zero profits, is $\mathcal{P} \equiv \{(P_{90}, P_{60}) : \Pi_{90}(P_{90}, P_{60}) = \Pi_{60}(P_{90}, P_{60}) = 0\}$. We also let $\underline{\Delta P}^{BE}$ denote the *lowest break-even* ΔP *with positive sales of the 60 policy.* This is the lowest price difference among all break-even price configurations with positive sales of the 60 policy, defined formally as:

$$\underline{\Delta P}^{BE} \equiv \min\{\Delta P : \text{there is a } (P_{90}, P_{60}) \in \mathcal{P} \text{ with } \Delta P = P_{90} - P_{60} > \underline{\theta}\}. \tag{1}$$

Note that the price configuration $(P_{90}, P_{60}) = (\overline{AC}_{60} + \overline{\theta}, \overline{AC}_{60})$, which results in all consumers purchasing the 60 policy, is a break-even price configuration (i.e., it is in set $\mathcal{P}$), as is the "all-in-90" price configuration $(P_{90}, P_{60}) = (\underline{AC}_{90}, \underline{AC}_{90} - \underline{\theta})$. There may also be "interior" break-even price configurations, at which both policies have a positive market share. The price difference $\underline{\Delta P}^{BE}$ will play a significant role in our equilibrium characterizations below.

2.1 Equilibrium Notions and Characterizations

The literature on equilibria in insurance markets with adverse selection started with Rothschild and Stiglitz (1976). Motivated by the possibility of non-existence of equilibrium in their model, follow-on work by Riley (1979) [see also Engers and Fernandez (1987)] and Wilson (1977) proposed alternative notions of equilibrium in which existence was assured in the Rothschild-Stiglitz model. These alternative equilibrium notions each incorporated some kind of dynamic reaction to deviations [introduction of additional profitable policies in Riley (1979), and dropping of unprofitable policies in Wilson (1977)], in contrast to the Nash assumption made by Rothschild and Stiglitz. In addition, follow-on work also allowed for multi-policy firms [Miyazaki (1977)], in contrast to Rothschild and Stiglitz's assumption that each firm offers at most one policy.

Our model differs from the Rothschild-Stiglitz setting in three basic ways. First, the prescription of health exchanges limits the set of allowed policies. Figure 1, for example, shows the set of feasible policies in the Rothschild-Stiglitz model (in which each consumer faces just two health states: "healthy" and "sick") with two exchanges, one for a 90% policy and the other for a 60% policy. These lie on lines with slope equal to 1 since a decrease of $1 in a policy's premium increases consumption by $1 in each state. Second, in our model consumers face many possible health states. Third, while the Rothschild-Stiglitz model contemplated just two consumer types, we assume there is a continuum of types of consumers.[14]

[14] To our knowledge, no existing work analyzes equilibria in insurance markets with these features. Einav and Finkelstein

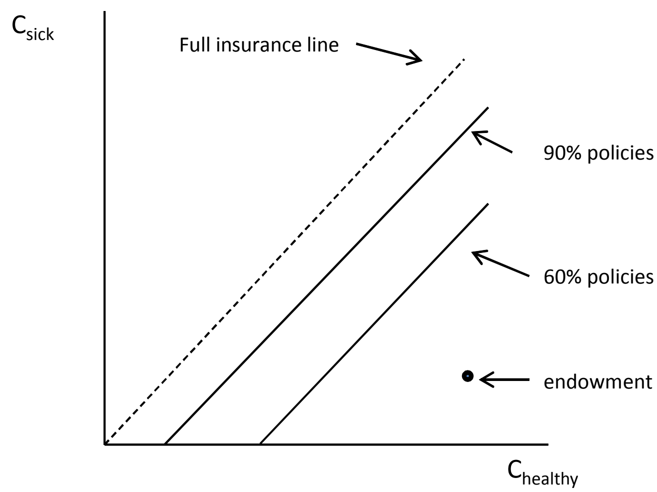

Figure 1: The solid lines with slope equal to 1 indicate the possible consumptions arising with 90% and 60% policies in a two-state (Rothschild-Stiglitz) model of insurance

In our main analysis we focus on the Riley equilibrium ("RE") notion, which we show always exists and is unique in our model.[15] We also discuss how these compare to Nash equilibria ("NE"), which need not exist. (In addition, we consider Wilson equilibria in Section 6.6.) In what follows, the phrase *equilibrium outcome* refers to the equilibrium price configuration and the shares of the two policies. Finally, to simplify the statement of the results, we restrict attention to equilibria with a price difference $\Delta P \in [\underline{\theta}, \overline{\theta}]$. Equilibria with $\Delta P < \underline{\theta}$ (resp. $\Delta P > \overline{\theta}$) exist if and only if one exists with $\Delta P = \underline{\theta}$ (resp. $\Delta P = \overline{\theta}$), and yield identical market shares, utilities, and profits.

2.1.1 Nash Equilibria

We consider Nash equilibria with both single-policy and multi-policy firms ("sp-NE" and "mp-NE", respectively). The following result characterizes these equilibria in our model (all proofs are contained in the Appendix):

Proposition 1. *With either single-policy or multi-policy firms, any NE price configuration* (P_{90}^*, P_{60}^*) *must have firms break even on all policies that are sold in equilibrium. If* $\Pi_{60}(\underline{AC}_{90}, P_{60}) \leq 0$ *for*

(2011) analyze a model with the latter two characteristics but just one policy type using a graphical price-theoretic approach.

[15] The Riley notion is also known as a "reactive equilibrium."

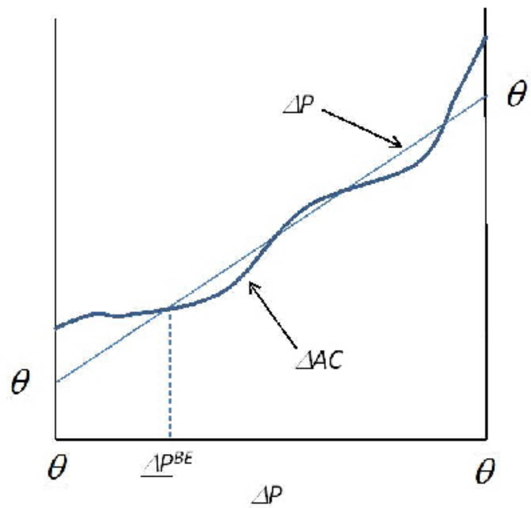

Figure 2: The figure shows $\underline{\Delta P}^{BE}$, the lowest price difference in any break-even price configuration that has positive sales of the 60 policy. It also shows a situation in which all-in-90 is not an equilibrium outcome, because $\Delta AC(\underline{\theta}) > \underline{\theta}$.

all P_{60} (i.e., if there is no profitable entry into the 60 policy given that the 90 policy is priced to break even), then the unique NE outcome has all consumers buying the 90 policy at price $P^*_{90} = \underline{AC}_{90}$. If this condition does not hold [which necessarily is the case when $\Delta AC(\underline{\theta}) > \underline{\theta}$], then any NE price configuration (P^*_{90}, P^*_{60}) must have price difference $\Delta P^* = \underline{\Delta P}^{BE}$, the lowest break-even ΔP with positive sales of the 60 policy. Such a price configuration (P^*_{90}, P^*_{60}) is a NE for:

(i) *single-policy firms if there is no profitable entry opportunity in the 90 policy; i.e., if $\Pi_{90}(\widehat{P}_{90}, P^*_{60}) \leq 0$ for all $\widehat{P}_{90} \leq P^*_{90}$;*

(ii) *multi-policy firms if there is no profitable entry opportunity that slightly undercuts P^*_{60} and undercuts P^*_{90}: i.e., if $\max_{\widehat{P}_{90} \leq P^*_{90}} \Pi(\widehat{P}_{90}, P^*_{60}) = 0$.*

The result says that all consumers buying the 90 policy can be a NE only if that outcome is immune from deviations that "cream skim," lowering P_{60} to attract the healthiest consumers to the 60 policy. If a cream-skimming deviation does break the all-in-90 outcome, then any NE must involve the price difference $\underline{\Delta P}^{BE}$. That price difference is illustrated in Figure 2, which plots $\Delta AC(\Delta P)$. The price differences at interior break-even price configurations are the $\Delta P \in (\underline{\theta}, \overline{\theta})$ at which the $\Delta AC(\Delta P)$ curve crosses the ΔP line. The figure also illustrates a situation in which $\Delta AC(\underline{\theta}) > \underline{\theta}$, implying that all-in-90 is not a Nash equilibrium.

The difference noted between single-policy and multi-policy deviations in parts (i) and (ii) of Proposition 1 arises because a price cut in P_{90} makes the 60 policy earn positive profits by attracting away its highest cost consumers. Thus, when an entrant can offer multiple policies it will want to slightly undercut P^*_{60} in order to retain the consumers who still buy the 60 policy.

2.1.2 Riley Equilibria

We use (a slightly modified version of) the definition provided in Engers and Fernandez (1987):

Definition 1. *A **Riley equilibrium** (RE) is a profitable market offering S, such that for any non-empty set S' (the deviation), where $S \cup S'$ is closed and $S \cap S' = \emptyset$, if S' is strictly profitable when $S \cup S'$ is offered then there exists a set S'' (the reaction), disjoint from $S \cup S'$ with $S \cup S' \cup S''$ closed, such that:*

(i) *S' incurs losses when $S \cup S' \cup S''$ is tendered;*

(ii) *S'' does not incur losses when any market offering $\widehat{S}$ containing $S \cup S' \cup S''$ is tendered (we then say S'' is "safe" or a "safe reaction").*

*A deviation S' that is strictly profitable when $S \cup S'$ is offered, and for which there is no safe reaction S'' that makes S' incur losses (with market offering $S \cup S' \cup S''$), is a **profitable Riley deviation**.*

In our setting, a market offering is simply a collection of prices offered for the two policies. Definition 1 says that a set of offered prices is a Riley equilibrium if no firm, including potential entrants, has a profitable deviation that also never leads it to incur losses should other firms introduce additional "safe" price offers (where a "safe" price offer is one that would never incur losses were any *further* price offers introduced).[16]

Our result for Riley equilibria is the following:

Proposition 2. *There is a unique Riley equilibrium. Moreover:*

(i) *If $\Pi_{60}(\underline{AC}_{90}, P_{60}) \leq 0$ for all P_{60} (i.e., if there is no profitable entry into the 60 policy given that the 90 policy is priced to break even), it involves everyone buying the 90 policy at price $P^*_{90} = \underline{AC}_{90}$.*

(ii) *Otherwise, it involves the break-even price configuration (P^*_{90}, P^*_{60}) with price difference $\Delta P^* = \underline{\Delta P}^{BE}$, the lowest break-even ΔP with positive sales of the 60 policy.*

Propositions 1 and 2 imply that all consumers buying the 90 policy is the unique equilibrium outcome at price $P^*_{90} = \underline{AC}_{90}$ under the exact same circumstances with both the NE and RE concepts. Where they differ is in what happens when this is not true. In both NE and RE any equilibrium must then involve price difference $\underline{\Delta P}^{BE}$, the lowest break-even ΔP with positive sales of the 60 policy.[17] However, under RE, this is always an equilibrium when all-in-90 is not an equilibrium. Under the two

[16] In fact, it suffices to restrict attention to deviations by potential entrants.

[17] Observe also that the outcome associated with $\underline{\Delta P}^{BE}$ then Pareto dominates the outcome associated with any other break-even price configuration that has positive sales of the 60 policy, as average costs, and hence prices, are lower.

NE concepts, however, an equilibrium may fail to exist, with the exact conditions for this depending on whether there are single-policy or multi-policy firms. The break-even price configuration with price difference $\underline{\Delta P}^{BE}$ can be a RE when it fails to be a NE because under the RE concept a profitable Nash deviation can be rendered unprofitable by additional profitable (and "safe") entry once the initial deviation occurs.

Recall that there are always break-even price differences that result in all consumers buying the 90 policy or all buying the 60 policy, while there may be (several) interior break-even price differences as well. Proposition 2 identifies the relevant one.

3 Data and Estimation

To simulate equilibria in health insurance exchanges we need a population of insurees, their preferences, and health status measures. This section describes the data that we use to obtain these ingredients, our empirical model, and the estimates. While the estimation is based on Handel (2013) we expand on that empirical model in several ways. Most importantly, we model consumer risk preference heterogeneity more flexibly by allowing for correlations with health risk, and include additional dimensions of observable heterogeneity, such as income and job type. These additional features are motivated by the empirical literature on adverse selection and insurance plan choice, which illustrates that correlations between risk preferences and risk can have important implications for equilibrium outcomes [see, e.g., Finkelstein and McGarry (2006), Cohen and Einav (2007), and Einav et al. (2013)].

3.1 Data

Our analysis uses detailed administrative data on the health insurance choices and medical utilization of employees (and their dependents) at a large U.S.-based firm over the time period from 2004 to 2009. These proprietary panel data include the health insurance options available in each year, employee plan choices, and detailed, claim-level employee (and dependent) medical expenditure and utilization information. While the employees at the firm are not 'representative' of any specific policy-relevant exchange population, the data are well-suited to estimate the ingredients necessary to illustrate equilibrium in exchanges. Later in the paper (Section 6) we perform an analysis that matches our sample to the nationally representative MEPS data, which we find is similar to our sample on a variety of dimensions and leads to similar results. We describe the data at a high-level in this section: for a more in-depth description of different dimensions see Handel (2013).

The first column of Table 1 describes the demographic profile of the 11,253 employees who work at the firm for some period of time within 2004-2009 (the firm employs approximately 9,000 at one time). These employees cover 9,710 dependents, implying a total of 20,963 covered lives. 46.7% of the employees are male and the mean employee age is 40.1 (median of 37). The table also presents statistics on sample income, family composition, and employment characteristics.

Sample Demographics	All Employees	PPO Ever	Final Sample
N - Employee Only	11,253	5,667	2,023
N - All Family Members	20,963	10,713	4,544
Mean Employee Age (Median)	40.1	40.0	42.3
	(37)	(37)	(44)
Gender (Male %)	46.7%	46.3%	46.7%
Income			
Tier 1 ($<$ \$41K)	33.9%	31.9%	19.0%
Tier 2 (\$41K-\$72K)	39.5%	39.7%	40.5%
Tier 3 (\$72K-\$124K)	17.9%	18.6%	25.0%
Tier 4 (\$124K-\$176K)	5.2%	5.4%	7.8%
Tier 5 ($>$ \$176K)	3.5%	4.4%	7.7%
Family Size			
1	58.0 %	56.1 %	41.3 %
2	16.9 %	18.8 %	22.3 %
3	11.0 %	11.0 %	14.1 %
4+	14.1 %	14.1 %	22.3 %
Staff Grouping			
Manager (%)	23.2%	25.1%	37.5%
White-Collar (%)	47.9%	47.5%	41.3%
Blue-Collar (%)	28.9%	27.3%	21.1%
Additional Demographics			
Quantitative Manager	12.8%	13.3%	20.7%
Job Tenure Mean Years (Median)	7.2	7.1	10.1
	(4)	(3)	(6)

Table 1: This table presents summary demographic statistics for the population we study. The first column describes demographics for the entire sample whether or not they ever enroll in insurance with the firm. The second column summarizes these variables for the sample of individuals who ever enroll in a PPO option, the choices we focus on in the empirical analysis. The third column describes our final estimation sample, which includes those employees who (i) are enrolled in PPO_{-1} at t_{-1} and (ii) remain enrolled in any plan at the firm through at least t_1.

Our analysis focuses on a three-year period in the data beginning with a year we denote t_0. For t_0, which is in the middle of our observational period, the firm substantially changed the menu of health plans that it offered to employees. At the time of this change, the firm forced all employees to leave their prior plan and actively re-enroll in one of five options from the new menu, with no default option. These five options were comprised of three PPO options, which shared the same broad provider network, and two HMO options, which led to some cost savings through different, narrower, provider networks. Our analysis focuses on choice among the three PPO options, which approximately 60% of health plan enrollees chose. We focus on this subset of the overall option set because (i) we have detailed claims data for PPO enrollees but not for HMO enrollees and (ii) the PPO options share the same doctors / cover the same treatments, eliminating a dimension of heterogeneity that would have to be identified separately from risk preferences. Analysis in Handel (2013) reveals, reassuringly, that while there is substitution across options within the set of PPO options, and across the set of HMO options, there is little substitution between these two subsets of plans, implying there is little loss of internal validity when considering choice between just the set of PPO options.

Within the nest of PPO options, consumers chose between three non-linear insurance contracts that differed on financial dimensions only. We denote the plans by their individual level deductibles: PPO_{250}, PPO_{500}, and PPO_{1200}. Post-deductible, the plans have coinsurance rates ranging from 10% to 20%, and out-of-pocket maximums after which the family spends no more out-of-pocket as total medical expenditures increase. PPO_{250} is the most comprehensive plan (i.e., provides the most financial protection) and thus has the lowest deductible, coinsurance, and out-of-pocket maximums (which depend on income as well as the number of dependents covered). PPO_{1200} is the least comprehensive plan on all financial dimensions. In terms of actuarial equivalence value (the proportion of expenditures covered for a representative population), PPO_{250} is approximately a 90% actuarial equivalence value plan while PPO_{1200} is approximately a 73% actuarial equivalence value plan (PPO_{500} is about halfway between PPO_{250} and PPO_{1200}). The plans have (subsidized) up-front premiums that are highest for PPO_{250}, lowest for PPO_{1200}, and depend on both income and the family members covered.[18] Over the three-year period that we study, t_0 to t_2, there is substantial variation in the premiums for these plans; this variation is helpful for identifying risk preferences separately from consumer inertia. For more details on the respective plan designs, and the evolution of premiums, see Handel (2013).

We restrict the final sample used in choice model estimation to those individuals / families that (i) enroll in one of the three PPO options and (ii) are present in all years from t_{-1}, the year before the menu change, through at least to t_1, one year before the end of our study period.[19] The reasons for the

[18] PPO_{1200} also has a linked Health Savings Account (HSA) option that allows consumers to deposit funds that can be used for medical expenditures on a pre-tax basis. This bundled account may be attractive due to tax-savings but unattractive because of increased hassle costs. We account for this feature in the estimated choice model. See Handel (2013) for a further discussion.

[19] We model plan choice in a given year based on the number of family members enrolled at the beginning of the year. We don't model changes to the number of dependents during a given year (or potential resulting changes to plan enrollment), since this occurs rarely and would complicate the analysis. For new dependents with no past health data

first restriction are discussed above. The second restriction, to more permanent employees, is made to leverage the panel nature of the data, especially the temporal variation in premiums and health risk, to more precisely identify risk preferences. Moreover, this more permanent population is simpler to model in the sense that their choices are always for the full year in advance and we always observe full past years of medical histories. Column 2 in Table 1 presents the summary statistics for the families that choose one of the PPO options, while Column 3 presents the summary statistics for the final estimation sample, incorporating the additional restriction of being present from t_{-1} to at least t_1. Comparing the second column to the first column reveals little selection on demographic dimensions into the PPO options, while comparing the third column to the others reveals some selection based on family size and age into the final sample, as expected given the restriction to longer tenure.

3.2 Health Status

We use detailed medical and demographic information together with the "ACG" software developed at Johns Hopkins Medical School to create individual-level measures of predicted expected medical expenses for the upcoming year at each point in time.[20] We denote these ex ante predictions of the next year's expected medical expenditures by λ and compute these measures for each individual in our observed sample (including dependents as well as employees). We refer to λ_{it} as individual i's "health status" at time t. We use these health status measures as inputs into our cost model, described in the next section and in Appendix B, to model uncertainty in health expenses for the upcoming year at the time of plan choice.

Health Status Descriptives

Figure 3 presents the distribution of λ for individuals in the data, as predicted for year t_1, for individuals (including dependents) present at both t_0 and t_1. The figure presents predicted health status (i.e., expected expenses) normalized by average predicted yearly expenditures of \$4,878 for these individuals for t_1. As is typical in the health care literature, the distribution is skewed with a large right tail (the chart truncates this right tail at 5 times the mean, though this is not done in our analysis). As we show later in Section 6, the distribution of expenditures in our population, both conditional and unconditional on age, is similar to that in the nationally representative MEPS survey data.

Table 2 describes health status transitions in the population over one and two year time horizons. This illustrates, from a short-run perspective, the potential for reclassification risk if premiums are allowed to depend, at least to some extent, on health status. The table studies transitions from year to year for quartiles of λ in the population: thus we see whether an individual transitions from one quartile of the health status distribution to another.[21] For this table, an observation for a one-year transition

we take the typical health expenditure distribution for someone of that age and gender.

[20] The program, known as the Johns Hopkins ACG (Adjusted Clinical Groups) Case-Mix System, is one of the most widely used and respected risk adjustment and predictive modeling packages in the health care sector. It was specifically designed to use diagnostic claims data to predict future medical expenditures.

[21] Note that this case of quartile transitions is directly relevant to the pricing case we study in the next sections, where

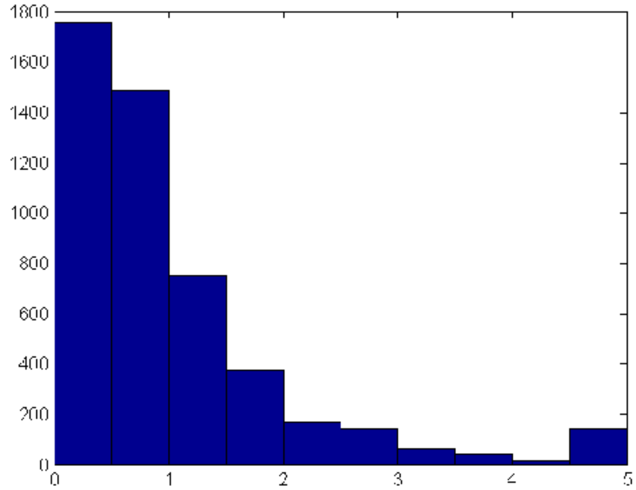

Figure 3: This figure presents the distribution of λ predicted for t_1, for all individuals in the data (including dependents) present during both t_0 and t_1. Predicted expected expenses are normalized by the average in the population of $4,878 (thus equal to 1 in this chart). The distribution presented is truncated at 5 times for this chart, but not in estimation / analysis.

is an individual in our data present over a given two-year time period (for the two-year transitions, it is any individual present over any given three-year period). Thus an individual present over four consecutive years will count as three observations in the one-year transition table.

The table reveals that there are real transition risks even for the fairly short one and two year time horizons: for example, 32% of the individuals in the healthiest quartile in year $t-1$ transition to one of the other three quartiles at year t (42% transition away from this quartile over a two-year period). To illustrate the potential for premium reclassification, the bottom section of the table presents average and median ex post cost by quartile grouping, indicating an increase in expected expenditures from $1,812 for quartile 1 to $15,199 for quartile 4. Note that since the table studies an aging population (not a steady state population) there is a trend towards higher health expenditures in these transitions.

3.3 Cost Model

The health status measure λ measures expected total health expenses. However, to evaluate the expected utility for consumers from different coverage options we need to estimate an ex ante distribution of out-of-pocket expenses for each family j choosing a given health plan k, not just their mean out-of-pocket expense. We utilize the cost model developed in Handel (2013) to estimate these distributions,

insurers are allowed to price based on health status quartile.

1 Year Transition

t-1 / t	λ Quartile 1	λ Q2	λ Q3	λ Q4
λ Quartile 1	0.68	0.18	0.08	0.06
λ Quartile 2	0.12	0.46	0.25	0.17
λ Quartile 3	0.05	0.15	0.39	0.41
λ Quartile 4	0.03	0.05	0.17	0.75

2 Year Transition

t-2 / t	λ Quartile 1	λ Q2	λ Q3	λ Q4
λ Quartile 1	0.58	0.22	0.09	0.11
λ Quatile 2	0.10	0.35	0.29	0.26
λ Quartile 3	0.03	0.14	0.28	0.56
λ Quartile 4	0.03	0.05	0.10	0.82

Cost Profile ($)

Quartile	Avg. Cost	Median Cost
λ Q1	1,812	302
λ Q2	3,544	1,107
λ Q3	5,543	2,542
λ Q4	15,199	6,831

Table 2: This table describes health status transitions in the population over one and two year time horizons. For the table, we group employees into ex ante health quartiles using λ. The top two sections describe these transitions, while the final section describes costs as a function of quartile.

denoted $H_k(X_{jt}|\boldsymbol{\lambda}_{\mathbf{jt}}, \mathbf{Z}_{jt})$.[22] Here, $\boldsymbol{\lambda}_{\mathbf{jt}}$ is the vector of λ_{it} for all i in family j, $\mathbf{Z}_{jt}$ are family demographics, and X_{jt} are out-of-pocket medical expenditure realizations for family j in plan k at time t. The model offers a parsimonious method to non-parametrically link health risk to expected future expenditures by combining the ACG software's predictive health risk measures with observed cost data.

We describe the details of the cost model in Appendix B; here we provide a broad overview of the methodology. The model has the following primary components:

1. For each individual and time period, we generate predictive mean expenditure measures for four categories of medical expenditures: (i) hospital/inpatient (ii) physician office visits (iii) mental health and (iv) pharmacy.

2. We next group individuals into cells based on mean predicted future utilization. For each expenditure type and risk cell, we estimate an expenditure distribution for the upcoming year based on ex post observed cost realizations. Then we combine the marginal distributions across expenditure categories into joint distributions using empirical correlations and copula methods.

3. Finally, we construct the detailed mappings from the vector of category-specific medical expenditures to plan out-of-pocket costs for each plan k. This includes plans in the actual data, to generate the distributions for choice model estimation, as well as candidate plans for the counterfactual simulations in Sections 4-6. For the plans in the observed data that we use for choice model estimation, this mapping inputs individual total expense projections and outputs family out-of-pocket expense projections taking into account family-level plan characteristics.

The output from this process, $H_k(X_{jt}|\boldsymbol{\lambda}_{\mathbf{jt}}, \mathbf{Z}_{jt})$, represents the distribution of out-of-pocket expenses associated with plan k used to compute expected utility in the choice model (and counterfactuals).

The cost model assumes both that there is no individual-level private information and no moral hazard (total expenditures do not vary with k). While both of these phenomena have the potential to be important in health care markets, and are studied extensively in other research, we believe that these assumptions do not materially impact our estimates. Because our cost model combines detailed individual-level prior medical utilization data with sophisticated medical diagnostic software there is less room for private information (and selection based on that information) than in prior work that uses coarser information to measure health risk.[23] For moral hazard, Chandra (2010) presents a recent review of the experimental and quasi-experimental literature, where the price elasticity for medical

[22] We note that the cost model leverages a multi-dimensional vector of health status measures, corresponding to risk for different types of medical expenditures (e.g., pharmacy, mental health, hospital). This allows for a richer / more precise model of H_{jkt}, and is described in detail in Appendix B.

[23] Pregnancies, genetic pre-dispositions, and non-coded disease severity are possible examples of private information that could still exist. Cardon and Hendel (2001) find no evidence of selection based on private information with coarser data while Carlin and Town (2009) use claims data that are similarly detailed to ours and also argue that significant residual selection is unlikely.

care generally falls in the range -0.1 to -0.4.[24] Online Appendix C in Handel (2013) illustrates that incorporating projected price elasticities for medical spending larger than these found in the literature has only a minor impact on the estimates of a choice model similar to ours.

3.4 Risk Preferences: Choice Model

We estimate risk preferences with a panel discrete choice model where choices are made by each household j at time t, conditional on their household-plan specific ex ante out-of-pocket cost distributions $H_k(X_{jt}|\boldsymbol{\lambda}_{\mathbf{jt}}, \mathbf{Z}_{jt})$. Specifically, the utility of plan k for household j at time t is:

$$U_{jkt} = \int_0^\infty u_j(W_j, X_{jt}, P_{jkt}, \mathbf{1}_{jk,t-1}, \mathbf{Z}_{jt}) H_k(X_{jt}|\boldsymbol{\lambda}_{\mathbf{jt}}, \mathbf{Z}_{jt}) dX_{jt} \tag{2}$$

Here, u_j is the v-NM or "Bernoulli" expected utility index that measures utility conditional on a given ex post realized state X_{jt} from the expenditure distribution H_k. W_j denotes income and P_{jkt} is the premium contribution for plan k at time t, which as described earlier depends both on how many dependents are covered and on employee income. $\mathbf{1}_{jk,t-1}$ is an indicator that equals one if plan k is the household's incumbent plan (default option) at choice year t. We use this variable to model consumer inertia, which is present for years with a default option (t_1 and t_2). Given that inertia is an important determinant of choices in those years, as shown in Handel (2013), we include this here to ensure that we appropriately identify risk preferences separately from this inertia. $\mathbf{Z}_{jt}$ are other individual-level observables (described shortly).

We assume that households have constant absolute risk aversion (CARA) preferences, leading to the utility index:

$$u_j(M_{jkt}) = -\frac{1}{\gamma_j(\mathbf{Z}_j^A, \boldsymbol{\lambda}_j)} e^{-\gamma_j(\mathbf{Z}_j^A, \boldsymbol{\lambda}_j) M_{jkt}} \tag{3}$$

Here, M_{jkt} is the effective consumption for a household given their ex post realization of health expenditures X_{jt} from distribution H_k and equals:

$$M_{jkt} = W_j - P_{jkt} - X_{jt} + \eta(\mathbf{Z}_j^B)\mathbf{1}_{jk,t-1} + \delta_j(A_j)\mathbf{1}_{1200} + \alpha HTC_{j,t-1}\mathbf{1}_{250} + \varepsilon_{jkt}(A_j) \tag{4}$$

Thus, consumption for a given health state realization equals household wealth, minus the up-front premium for plan k at time t, minus the out-of-pocket health expenditures X_{jt}. In addition, we model inertia, as a function of observable heterogeneity in demographics $\mathbf{Z}_j^B$, similarly to a tangible switching cost: consumption for every ex post state X_{jt} is the equivalent of $\$\eta$ lower if the plan chosen is not the same as the default option (i.e., the consumer has to incur cost η to switch).[25] $\delta_j(A_j)$ is a random coefficient, with distribution estimated conditional on family status A_j (single or covering dependents),

[24] Recent work by Einav et al. (2013), with data from a large employer, finds an implied elasticity of -0.14. The well-known RAND experiment has an oft-cited elasticity estimate of -0.18.

[25] See Handel (2013) for an extended discussion of this model for inertia, and how it relates to other potential microfoundations for inertia (i.e., other than a tangible switching cost / transaction cost).

that captures permanent horizontal preferences for PPO_{1200} arising from the Health Savings Account linked to this plan option. α captures preferences for very high-expenditure consumers, who almost exclusively choose PPO_{250} even when that option is not attractive financially ($HTC_{j,t-1} = 1$ for the top 10% of the distribution of expected total costs).[26] The utility of each option k for family j at t is also affected by a mean zero idiosyncratic preference shock ε_{jkt} known to the decision-maker, with variance σ_ε to be estimated conditional on A_j.

γ_j is a household-specific CARA risk preference parameter unobserved by the econometrician that depends on observed demographics $\mathbf{Z}_j^A$ and expected household expenditures (the sum of λ_i over $\boldsymbol{\lambda}_j$). γ_j determines the curvature of $u_j(M_{jkt})$ and, consequently, household risk aversion with respect to the lottery over consumption M_{jkt} induced by the distribution of out-of-pocket health expenditures H_k. We estimate a random-coefficient distribution of γ_j that is assumed to have mean $\mu_\gamma(\mathbf{Z}_j^A, \boldsymbol{\lambda}_j)$ and be normally distributed variance σ_γ^2. Note that observable heterogeneity impacts risk preference estimates through a shift in μ_γ, while the level of unobserved heterogeneity measured by σ_γ^2 is assumed constant for the entire population. We use the following specification for $\mu_\gamma(\mathbf{Z}_j^A, \boldsymbol{\lambda}_j)$:

$$\mu_\gamma(\mathbf{Z}_j^A, \boldsymbol{\lambda}_j) = \beta_0 + \beta_1 \log(\Sigma_{i\epsilon j}\lambda_i) + \beta_2 age_j + \beta_3 \log(\Sigma_{i\epsilon j}\lambda_i) * age_j + \beta_4 1_{mj} + \beta_5 1_{mj}\widehat{\upsilon}_{mj} + \beta_6 1_{nmj}\widehat{\nu}_{nmj} \quad (5)$$

In addition to expected household health expenditures ($\Sigma_{i\epsilon j}\lambda_i$), risk preferences depend on maximum household age, denoted age_j, and the interaction between health risk and age. 1_{mj} is an indicator variable that denotes whether the employee associated with the household is a "manager" (i.e., a high-level employee) at the firm. 1_{nmj} is the complement of 1_{mj}. $\widehat{\upsilon}_{mj}$ is a measure of ability, and is computed as the residual to the following regression, run only on the sample of managers in the population:

$$Income_{jt} = \alpha_0 + \alpha_1 age_{jt} + \alpha_2 age_{jt}^2 + \upsilon_{jt} \quad (6)$$

The residual $\widehat{\upsilon}_{nmj}$ is computed from the corresponding regression for non-managers.

Regarding identification, risk preferences are identified separately from inertia by leveraging the firm's insurance menu re-design for year t_0. Households in that year chose plans from a new menu of options with no default option, while in subsequent years they did have their previously chosen option as a default option. Conditional on this choice environment, changing prices and health status over time separately identify inertia from risk preference levels and risk preference heterogeneity. The different components of risk preference heterogeneity are identified by using exogenous price differences across both income tiers and coverage tiers (number of family members covered) and over time, as well as

[26] See Handel (2013) for a further discussion of both the Health Savings Account feature, linked only to PPO_{1200} (a "high-deductible" health plan), as well as the choice patterns of high-cost consumers. That paper discusses the potential for horizontal differentiation caused by the Health Savings account, and illustrates the importance of modeling this given observed choices and costs. In addition, it shows that, even in cases where PPO_{250} is dominated by the other two plan options, very high-cost consumers choose that plan with high frequency, perhaps because of a heuristic to choose the most comprehensive plan available. We abstract from this effect in our counterfactual simulations.

changes to household expenditure distributions over time. Prices change substantially across income tiers and family tiers, while across these tiers households can have similar expenditure risk distributions. Changes over time in health status and premiums, assuming risk preferences are constant over time, also provide identifying variation for risk preferences. Finally, consumer preference heterogeneity for the high-deductible plan option with the linked health savings account (HSA) is distinguished from risk preference heterogeneity by comparing choices between the two other plans to those between either of those plans and the high-deductible plan.

We estimate the choice model using a random coefficients simulated maximum likelihood approach similar both to that summarized in Train (2009) and to that used in Handel (2013). The likelihood function at the household level is computed for a sequence of choices from t_0 to t_2, since inertia implies that the likelihood of a choice made in the current period depends on the previous choice. Since the estimation algorithm is similar to a standard approach, we describe the remainder of the details, including the specification for heterogeneity in inertia, in Appendix C.

3.5 Preference Estimates

Table 3 presents our choice model estimates. The first column presents the estimates of our primary specification while the second through fourth columns present robustness analyses to assess the impact of linking different types of observable heterogeneity to risk preferences. The table presents detailed risk preference estimates, including the links to observable and unobservable heterogeneity. Since we don't use any other estimated parameters in the upcoming exchange equilibrium analyses (except for σ_ε), for simplicity we present and discuss the rest of the estimated parameters in Appendix C (e.g., inertia estimates, PPO_{1200} random coefficients, ε_{jkt} standard deviations, and income regressions). Parameter standard errors, which are generally quite small, are also presented in Appendix C.

For the primary specification, the population mean for μ_γ, the household mean risk-aversion level given unobserved heterogeneity, is $4.39*10^{-4}$. The standard deviation for μ_γ (or the standard deviation in risk preferences based on observable heterogeneity) equals $6.63*10^{-5}$. σ_γ (the standard deviation of unobservable heterogeneity in risk preferences) equals $1.24*10^{-4}$. In terms of observable heterogeneity, risk preferences are negatively correlated with health risk: a one point increase in $\log(\lambda)$ reduces μ_γ by $8.10*10^{-5}$ for a 30-year old.[27] This suggests that there should be less adverse selection in the simulated markets that we study relative to the case in which risk preferences are independent of health risk. Managers and those with higher ability are slightly more risk averse. With a log expected total health spending value of 9 (around the median for a household) risk aversion is increasing in age by $4.69*10^{-6}$ per year. The specifications in the second through fourth columns in the table, which investigate robustness with respect to the inclusion of and specification for health status / income

[27] The coefficient on health risk is more negative than this, while the interaction between age and risk preferences has a positive coefficient, indicating some reduction in the negative relationship between risk preferences and health risk as one becomes older.

Empircal Model Results				
	(1)	(2)	(3)	(4)
Parameter / Model	Primary Model	Robustness 1	Robustness 2	Robustness 3
Risk Preference Estimates				
μ_γ - Intercept, β_0	$1.21 * 10^{-3}$	$1.63 * 10^{-4}$	$1.06 * 10^{-3}$	$2.54 * 10^{-4}$
μ_γ - $\log(\Sigma_{i \epsilon j} \lambda_i)$, β_1	$-1.14 * 10^{-4}$	-	-$1.21 * 10^{-4}$	-
μ_γ - age, β_2	$-5.21 * 10^{-6}$	$3.60 * 10^{-6}$	-$4.69 * 10^{-6}$	$3.99 * 10^{-6}$
μ_γ - $\log(\Sigma_{i \epsilon j} \lambda_i)$$*$age, β_3	$1.10 * 10^{-6}$	-	$1.01 * 10^{-6}$	-
μ_γ - Manager, β_4	$4.3 * 10^{-5}$	$7.45 * 10^{-5}$	$5.3 * 10^{-5}$	$5.4 * 10^{-5}$
μ_γ - Manager ability, β_5	$1.4 * 10^{-5}$	$4.49 * 10^{-5}$	$-$	-
μ_γ - Non-manager ability , β_6	$7.5 * 10^{-6}$	$3.24 * 10^{-5}$	$-$	-
μ_γ - Nominal Income, β_7	-	-	$3.0 * 10^{-5}$	-
μ_γ - Population Mean	$4.39 * 10^{-4}$	$3.71 * 10^{-4}$	$4.33 * 10^{-4}$	$4.73 * 10^{-4}$
μ_γ - Population σ	$6.63 * 10^{-5}$	$7.45 * 10^{-5}$	$8.27 * 10^{-5}$	$6.30 * 10^{-5}$
σ_γ - γ standard deviation	$1.24 * 10^{-4}$	$1.14 * 10^{-4}$	$1.40 * 10^{-4}$	$1.20 * 10^{-4}$
Gamble Interp.:				
μ_γ Mean	693	728	696	676
μ_γ Mean + 25th Quantile σ_γ	736	772	748	717
μ_γ Mean + 75th Quantile σ_γ	653	688	651	640
μ_γ Mean + 95th Quantile σ_γ	604	638	596	593

Table 3: This table presents the our choice model estimates. The first column presents the results from our primary specification described in Section 3. The second through fourth columns present robustness analyses that assess the impact of linking preferences to health status and our measure of income earning ability. For each model, we present the detailed risk preference estimates, including the links to observable and unobservable heterogeneity. The rest of the parameters (inertia estimates, PPO_{1200} random coefficients, and ε_{jkt} standard errors) are provided in Appendix C. The bottom of the table interprets the population mean risk preference estimates: it provides the value X that would make someone indifferent about accepting a 50-50 gamble where you win \$1000 and lose X versus a status quo where nothing happens. The population distributions of risk preferences are similar across the specifications, even though the additional links between health risk / income and risk preferences add richness.

in risk preferences, estimate similar means and variances for risk preferences relative to our primary specification.

The bottom rows in Table 3 interpret the mean of the average estimated risk aversion μ_γ, as well as several quantiles surrounding that average μ_γ. We present the value X that would make a household with our candidate risk aversion estimate indifferent between inaction and accepting a simple hypothetical gamble with a 50% chance of gaining \$1000 and a 50% chance of losing \$$X$. Thus, a risk neutral individual will have $X = \$1000$ while an infinitely risk averse individual will have X close to zero. For the population mean of μ_γ from the primary model we have $X = \$693$ while for the 25th, 75th, and 95th quantiles of unobserved heterogeneity around that mean X is \$736, \$653 and \$604 respectively (these values are decreasing because they decrease as γ increases). While the estimates in the literature span a wide range, and should be interpreted differently depending on the different contexts being studied, our estimates generally fall in the middle of the range of prior work on insurance choice, while the extent of heterogeneity we estimate is somewhat lower in magnitude [see, e.g., Cohen and Einav (2007)]. Interestingly, the negative estimated correlation between expected health risk and risk preferences is consistent with that association in Finkelstein and McGarry (2006) but the opposite sign of the effect found in Cohen and Einav (2007).

3.6 Simulation Sample

For the choice model we estimate, it is necessary to estimate choice at the family level because that is the unit that actually makes choices in the data. For our counterfactual insurance exchange simulations, we focus on individuals to simplify exposition. In actual insurance exchanges this could be an appropriate model if family costs and premiums are aggregated from the individual level and family members are not required to enroll in the same plan. We note that the framework could easily be extended to allow for separate markets for individuals and families covering dependents, as is typical in practice.

The sample used in the simulations contains individuals between the ages of 25 and 65 who are present in our data. Thus, our simulations include both individuals with single coverage in the data, and individuals who are members of families with family coverage in our data. For the purposes of our simulation, individuals with family coverage in the data choose any individual coverage plan. Since we simulate static equilibria, but observe multiple years of health status and demographic variable realizations for individuals in our data, we include a given individual in our data multiple times in our simulation sample, based on their tenure in the actual data. For example, if an individual is present in our data for three full years, for example from ages 26 to 28, they are included as three separate individuals for the purposes of our simulation: a 26 year old individual, a 27 year old individual, and a 28 year old individual. Each of these simulated individuals has potentially different predicted health status, income, etc. based on their actual data for the relevant year in question. To ensure that the data for a given individual are complete, we require a given simulated individual to be present for at least eight months in each of two consecutive years.[28] The data from the first year are used to predict health status while the presence in the second year is used to ensure the individual was a relevant potential participant in the firm's benefit program for that year. This ensures that the simulation sample reflects to some extent the presence / longevity of the choice model estimation sample. For risk preferences, some of the variables used in estimation are defined at the family level rather than the individual level (e.g., income, manager status of the employee in the family). Every individual that comes from a given family is assigned the relevant family value for these variables when simulating risk preferences for that individual in the exchange counterfactuals.

Table 4 describes some key descriptive numbers for this pseudo-sample of 10,372 individuals used for the insurance exchange simulations. Importantly, the distribution of risk preferences in this sample is very similar to that in the estimation sample, implying it is not highly selected on this dimension. Similarly, the distributions of income and health expenditures are similar to those of the main estimation sample and the population overall. The proportion female is also similar. Finally, as shown below, the simulation sample covers the range of ages from 25-65 fairly evenly, which is reflective of this characteristic in our data in general. This is relevant to our upcoming welfare analysis, which assumes that the population is in a steady state.

Quantile	0.05	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9	0.95
Age	26	28	33	37	41	45	49	52	56	60	62

We now turn to our primary analysis of insurance exchange market equilibrium and welfare as a function of different underlying regulations.

4 Results: Equilibria

We use the estimates from our choice and cost models to study the effects of regulations that restrict the information insurers can use to price their policies.[29] As noted in Section 2, we mimic the exchange structure set forth in the ACA (see, e.g., Kaiser (2010)) and study a market in which insurers can offer one of two policies that cover either 90% or 60% of expenditures in the population, on average (recall, this is a simplification of the four classes of policies allowed under the ACA of 60%, 70%, 80%, and 90% average coverage). While there are a variety of potential non-linear contract designs that would imply these coverage levels, here we follow the discussion of such policies in Consumers Union (2009) and

[28] For individuals whose past year of cost data is less than one year (between eight months and one year) we assume that this past data represents one full year of health claims for the purposes of constructing their health status λ. We assume in all of the simulations that individuals buy a plan expecting to be in that plan for the full year (this is not an issue in choice model estimation, where the sample is restricted to those present for full years). The cost model estimation is done only for individuals with full years of cost data and these full-year distributions are the ones used in our analysis.

[29] We could also use our framework to investigate trade-offs the regulator faces when choosing the permissible contract space (in terms of plan cost-sharing features). One advantage of studying pricing regulation is that the contracts we observe in our data are similar to those used in our equilibrium simulations, which might not generally be the case if studying a wide range of contract design options.

Simulation Sample	Simulation Sample
N - Families	-
N - Individuals 25-65	10,372
Mean Age	44.5
Median Age	45
Gender (Male %)	45
Income	
Tier 1 ($<$ \$41K)	20%
Tier 2 (\$41K-\$72K)	40%
Tier 3 (\$72K-\$124K)	24%
Tier 4 (\$124K-\$176K)	8%
Tier 5 ($>$ \$176K)	8%
Predicted Mean Total Expenditures	
Mean	\$6,559
25th quantile	\$1,673
Median	\$3,675
75th quantile	\$8,354
90th quantile	\$13,937
95th quantile	\$18,638
99th quantile	\$33,835
Risk Preferences	
Mean μ_γ	$4.28 * 10^{-4}$
Standard Deviation μ_γ	$7.50 * 10^{-5}$

Table 4: This table presents descriptive statistics for the pseudo-sample of individuals used in our insurance exchange simulations. The sample has risk preference means and standard deviations that are similar to those of the choice model estimation sample. Moreover, the distributions of income and health status are similar to those in the estimation sample and general population.

assume that the 90% policy has no deductible, a 20% coinsurance rate post-deductible, and a $1500 out-of-pocket maximum (all at the individual level we study here) while the 60% policy has a $3,000 deductible, a 20% coinsurance rate post-deductible, and a $5,950 out-of-pocket maximum. These two plans cover 91% and 62% of expenditures for our entire simulation population on average, suggesting that these are appropriate plans for our analysis and that our population is similar on average to the representative population considered in this policy document.[30]

The estimated model contains three sources of heterogeneity that we use in this analysis: risk type, risk aversion, and an idiosyncratic preference shock. For each individual in the population we compute, based on their demographics and prior diagnostics, the risk type λ discussed in the previous section. Given that λ, we take 100 draws from the estimated distribution of γ conditional on λ and the other demographics modeled in equation (5), creating 100 "pseudo-individuals" for each actual individual in our sample. Doing so for each individual in the sample generates a joint distribution of risk preferences and risk type. For each of the two plan designs we study, we adapt the cost model described in Section 3 to estimate the person-plan specific distribution of out-of-pocket expenses $H_k(\cdot|\lambda_{it}, Z_{it})$ where individuals are indexed by i and the two plans by k.[31] With these objects, we compute the expected utility of each (pseudo) individual for each plan, and use them to find CE_{90} and CE_{60} (gross of premiums), as described in Section 2. Willingness to pay for the extra coverage of the 90% plan is $\theta = CE_{90} - CE_{60} + \varepsilon$, where ε is distributed $N(0, \sigma_\varepsilon^2)$. Thus, as in equation (4), there is a random shock to a consumer's preference between the two plans. For the simulations that follow we use $\sigma_\varepsilon = 525$, which is the estimated standard deviation of ε for the single population for PPO_{1200} relative to PPO_{250}.[32] As we report below, our results are robust to medium-sized changes in σ_ε.

The sample population and the estimated distributions determine $F(\theta)$. Costs to each plan k, $C_k(\theta)$ for $k = 90$ and 60, are computed using expected plan costs $\lambda_{it} - E[H_k(\cdot|\lambda_{it}, Z_{it})]$, aggregating over all individuals associated with each θ, while $AC_{90}(\theta)$ and $AC_{60}(\theta)$ are determined by aggregating these costs over the θ that select a given plan.

The Adverse Selection Property introduced in Section 2, upon which our theoretical results hinge, can be verified in our sample: Figure 4 shows that AC_{90} and AC_{60} are increasing in ΔP for each policy, and that AC_{90} exceeds AC_{60} at all ΔP.

[30] We also note that states can regulate the space of permissible non-linear contract designs for a given actuarially equivalence value (e.g., 90%). In practice, many states have chosen to allow insurers only a very restricted space of contracts within each class, making our assumption of one contract not very restrictive relative to reality. See, e.g., Ericson and Starc (2013) for a discussion of this regulation in Massachusetts.

[31] Recall, as described in Section 3, for these analyses we use individuals, not families, hence the notation i for individual rather than j for family as in the choice model.

[32] We choose these two plans because they are closest in our observed plan set to the 90% and 60% plans we study in our exchange analysis.

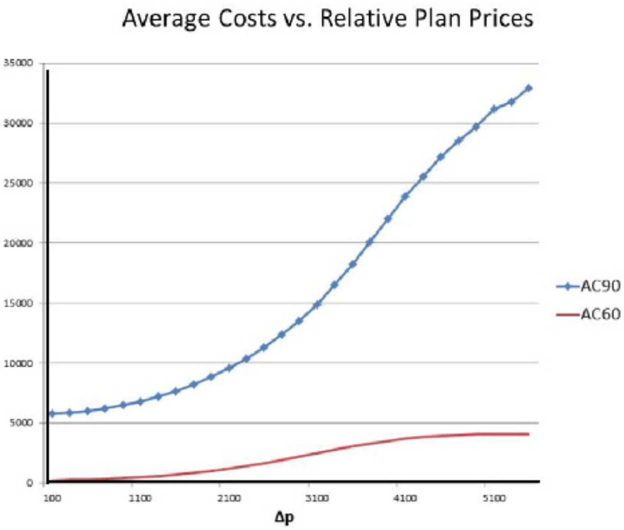

Figure 4: Plot of average costs vs. the price difference ΔP. Average costs are increasing in this price difference, and are larger for the 90 policy at each ΔP, consistent with the assumption maintained to derive our theoretical results.

4.1 No Pricing of Pre-Existing Conditions

We start by considering the case of pure community rating, where insurers must price everyone in the whole population identically. We follow the theoretical results of Section 2 as a roadmap to finding equilibria.

The first step towards finding equilibria involves checking whether all consumers pooling in the 90 plan, the highest level of coverage, is an equilibrium. For a 90 policy to break even covering all of the population, the premium P_{90} must equal $\underline{AC}_{90}$. Such a policy is an equilibrium if $\Pi_{60}(\underline{AC}_{90}, P_{60}) \leq 0$ for all P_{60}. If that inequality holds, Propositions 1 and 2 guarantee that all-in-90 is both a Nash and a Riley equilibrium, and that equilibrium is unique.

Figure 5, which plots $\Delta AC(\Delta P)$, shows that $\Delta AC(\underline{\theta}) > \underline{\theta}$ which guarantees that there is a profitable 60 deviation from all-in-90 by targeting the healthiest customers. Thus, in our population all-in-90 is not an equilibrium and the equilibrium must involve purchases of the 60 policy.

The second step towards finding either Riley or Nash equilibria involves finding the lowest break-even ΔP, $\underline{\Delta P}^{BE}$; i.e., the lowest interior ΔP at which $\Delta P = AC_{90}(\Delta P) - AC_{60}(\Delta P)$, if any exist, or $\Delta P = \overline{\theta}$ otherwise. This is then the RE ΔP, and is the only candidate for NE, should a NE exist.

Figure 5 shows that, for the case of pure community rating, there is no interior equilibrium. Namely, there is no pair of premiums at which both policies have positive market shares and both break even: for any premium gap between 60 and 90 coverage, the gap in costs due to adverse selection into 90 is larger than the gap in premiums. The market must fully unravel.

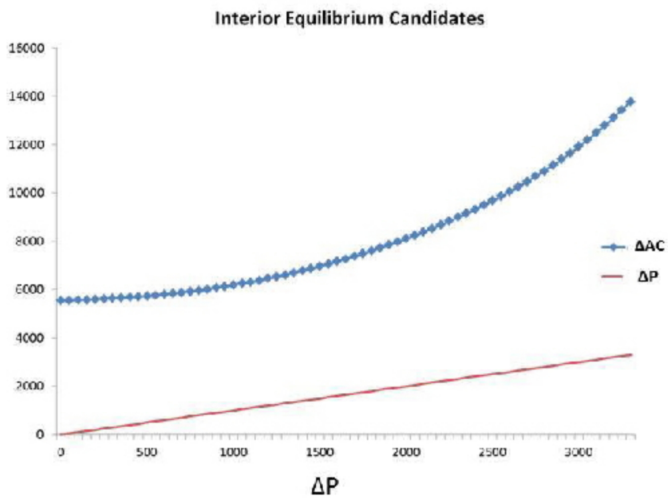

Figure 5: Plot of plan average costs as a function of their price differences.

The third, and final, step involves checking whether all-in-60 is an NE (Proposition 2 guarantees the existence of an RE, so having ruled out all candidates with positive 90 sales we know all-in-60 must be the RE). Unlike RE, the existence of an NE is not guaranteed.

Figure 6 shows the profits from both single-policy and multi-policy deviations starting with all consumers buying the 60 policy at the price $P_{60} = \overline{AC}_{60}$, namely, $\Pi_{90}(\overline{AC}_{60}+\Delta P, \overline{AC}_{60})$ and $\Pi(\overline{AC}_{60}+\Delta P, \overline{AC}_{60} - \varepsilon)$ for $\Delta P \leq \overline{\theta}$. As the lower curve shows, $\Pi_{90}(\overline{AC}_{60} + \Delta P, \overline{AC}_{60})$ is never positive. The worst risks, attracted into the 90 plan, are more costly than the premium they pay. Thus, pooling in the 60 plan is an sp-NE as well as an RE.

The higher total profits curve shows that a multi-policy deviation from all-in-60 is profitable. For such a deviation, while the 90 customers are not profitable by themselves, the pool left in the 60 plan, which can be attracted with $P_{60} = \overline{AC}_{60} - \varepsilon$, more than compensates for the losses on consumers shifting to the 90 plan.

The top section of Table 5 summarizes these findings for the case of a pure community rating pricing regulation.

4.2 Pricing Pre-existing Conditions

We now investigate the effects of allowing pricing of some health status information. Specifically, we consider the case in which consumers are classified into quartiles based on their ex ante predicted total expenditures λ : e.g., the first quartile contains all of the healthiest consumers, while the last contains all of the sickest consumers. Insurers can target each quartile with different prices as they see fit. We later present results that vary the fineness of information insurers can price on, ranging from pure

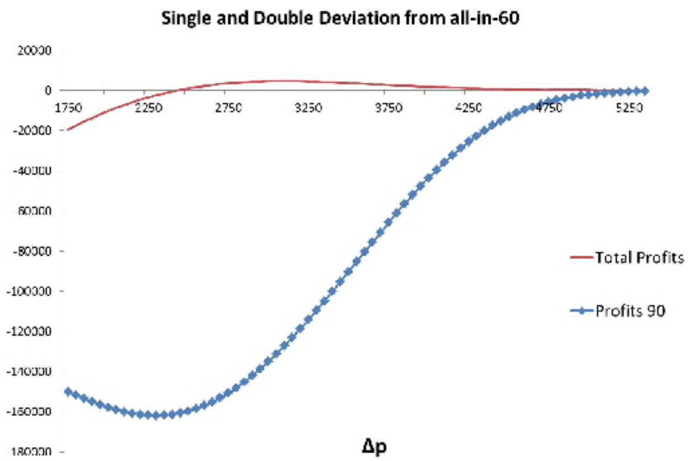

Figure 6: Investigation of pricing deviations from "All-In 60" candidate equilibrium.

Equilibria without Pre-existing Conditions						
Equilirium Type	P_{60}	Sh_{60}	AC_{60}	P_{90}	Sh_{90}	AC_{90}
Riley	4,051	100.0	4,051	–	0	–
Single-policy Nash	4,051	100.0	4,051	–	0	–
Multiple-policy Nash			Does not exist			

Equilibria with Health Status-based Pricing							
Market	Equilibrium Type	P_{60}	Sh_{60}	AC_{60}	P_{90}	Sh_{90}	AC_{90}
Quartile 1	RE/sp-NE/mp-NE	289	64.8	289	1,550	35.2	1,550
Quartile 2	RE/sp-NE	1,467	100.0	1,467	-	0	-
Quartile 3	RE/sp-NE	4,577	100.0	4,577	-	0	-
Quartile 4	RE/sp-NE	9,802	100.0	9,802	-	0	-

Table 5: The top section of this table presents the equilibrium results for the case of pure community rating (no pricing of pre-existing conditions). The bottom section presents the equilibrium results for the case where insurers can price based on health status information in the form of health status quartiles. The equilibrium results are presented for each health status quartile, which act as separate markets under this regulation.

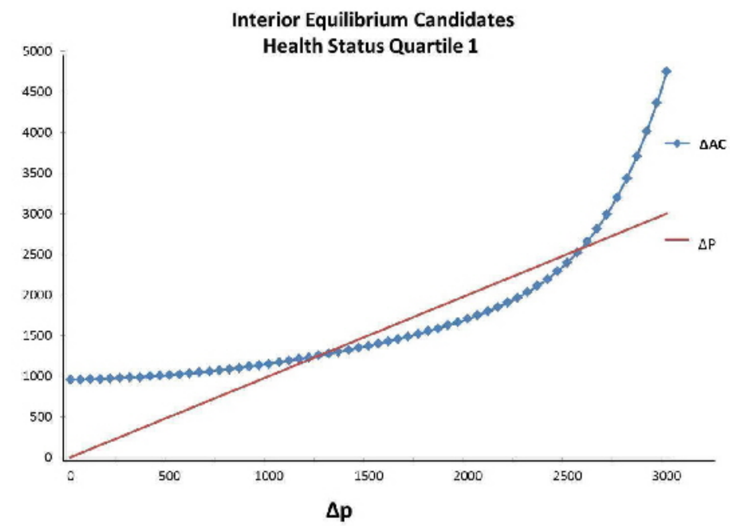

Figure 7: Potential deviations from candidate equilibrium where all consumers pool in the 90 plan.

community rating all the way up to the case of unrestricted risk rating / price discrimination. These alternative regulations are meant to be illustrative of potentially more subtle regulations seen in real-world insurance markets that increase the ability of insurers to price discriminate (e.g., pricing based on specific pre-existing medical conditions). We follow the same steps as in the previous subsection to find equilibria, but now for each market segment separately.

The implications of this pricing regulation for adverse selection are seen directly when examining the pricing equilibrium for quartile 1, the healthiest quartile of consumers. For quartile 1, there is an interior NE that survives multi-policy deviations and, thus, coincides with the RE. The first step, as described above, is to check whether all-in-90 is an equilibrium. Figure 7 shows that, as in the pure community rating case, $\Delta AC(\underline{\theta}) > \underline{\theta}$, implying that all-in-90 is not an equilibrium.

The second step is to look for interior equilibrium candidates. Figure 7 shows two interior break-even ΔPs. By Propositions 1 and 2 only the lowest ΔP, the one with the largest share of customers in the 90 plan is the RE and is the only candidate for a NE. Figure 8 displays the profitability of single and double deviations for the equilibrium associated with the lower break-even ΔP. As neither single-policy nor multi-policy deviations from this candidate ΔP are profitable, this ΔP is both a single- and multiple-policy NE.

In contrast, equilibria in quartiles 2, 3 and 4 are qualitatively identical to the equilibrium under pure community rating. The sp-NE is all-in-60 and coincides with the RE. An mp-NE does not exist. We omit the graphs, which look similar to Figures 5 and 6. The bottom section of Table 5 summarizes the findings for the four quartiles under health status-based pricing. The table also highlights the potential for reclassification risk when moving from the static equilibrium analysis to the analysis of long-run consumer welfare: if insurers can price based on health status quartiles, buyers will find themselves

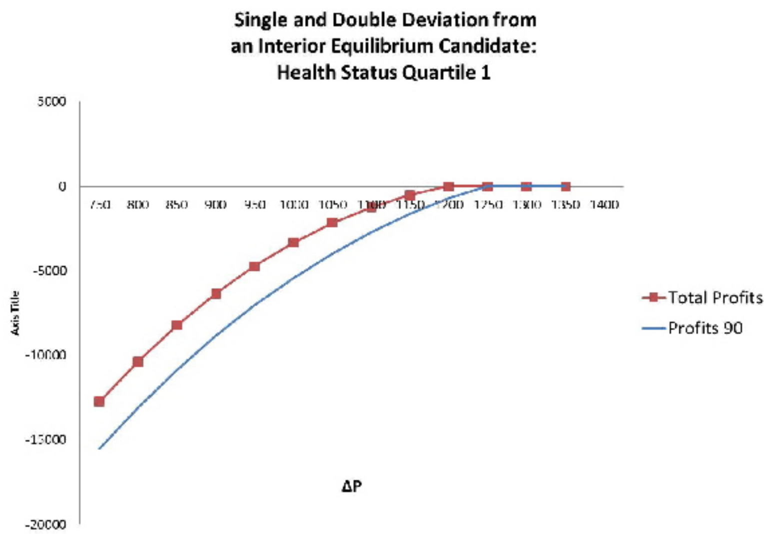

Figure 8: Profitability of deviations from interior equilibrium candidate with health status pricing.

paying premiums as low as \$289 or as high as \$9,802, corresponding to the different quartiles, as their health evolves over time. However, under these pricing regulations, many of the healthiest consumers in the population obtain a greater level of insurance coverage, and thus are less impacted by adverse selection.

To more completely analyze the trade-off between adverse selection and re-classification risk, we next consider a range of pricing regulations that allow insurers to price based on health status information with varying degrees of specificity. The second column in Table 6 describes the RE/sp-NE share in the 60 policy when insurers instead can price based on 2, 4, 6, 8, 10, 20, or 50 health status partitions, as well as the case of full risk-rating (labeled ∞). Adverse selection is reduced as the insurers are able to price on finer information: with 4, 10, and 50 partitions the 60 plan has 90%, 83%, and 63% market shares respectively, while with full risk-rating 73% of consumers choose to enroll in the 90 plan.[33] (The welfare numbers in columns 3-5 of Table 6 will be discussed in Section 5.)

5 Results: Welfare

Our aim in this section is to evaluate the expected utility of an individual starting at age 25 from an ex-ante ("unborn") perspective; that is, before he knows the evolution of his health. The unborn individual faces uncertainty about how his health status will transition from one year to the next, and thus what policies he will purchase and what premiums he will pay. Since individuals differ in their risk aversion, we will calculate this expected utility separately for different risk aversion levels.

[33] With no ε preference shock, with full risk-rating all consumers would enroll in the 90% plan. Here, with the estimated ε standard deviaiton incorporated, the first-best allocation has 73% of consumers in the 90% plan, since some prefer the 60% plan due to this preference shock.

Equilibria Welfare Loss from Health Status-based Pricing: Varying Regulation				
		$y_{x,no-pre}$	$y_{x,no-pre}$	$y_{x,no-pre}$
# of ACG Groups	Sh_{60}	Fixed Income	Non-Manager Income path	Manager Income Path
2	100.0	1,920	710	-102
4	90.0	3,082	1,821	-886
6	82.0	3,951	2,377	-232
8	85.1	4,649	2,084	-1,510
10	83.2	5,357	2,269	-1,364
20	81.4	8.590	4,621	-393
50	63.2	11,578	7,302	2,359
∞	27.0	14,733	9,944	2,399

Table 6: Equilibria and long-run welfare comparison between the pricing regulations that allow some pricing based on health status and the case where no pricing on health status is allowed. The table shows the share of consumers choosing the 60 policy for each pricing regime. It also presents the values for $y_{x,no-pre}$, the annual payment required under regime x to make consumers indifferent between x and no risk-rating. The regimes x listed in column 1 correspond to how targeted pricing can be over the range of health status: e.g., 4 corresponds to the case of quartile pricing while ∞ is full risk rating. The results presented are for Riley Equilibria and $\gamma = 0.0004$.

To be more specific, suppose we have pricing rule x (e.g., no pre-existing conditions) and equilibrium concept e (e.g., Riley equilibrium). The analysis in the previous section tells us what policy each individual will choose as a function of their health status (λ) and risk aversion (γ), and the premium they will pay. Given this information, we can compute the certainty equivalent $CE_{x,e}(\lambda, \gamma)$ of the uncertain consumption that this individual of type (λ, γ) will face within a year because of uncertainty over his health realization.

To compute expected utility starting at age 25 from an ex ante perspective, we need to know how health status will transition over time for an individual with a given risk aversion γ at age 25. To do so, we assume that our sample represents a steady state population in which the distribution of health types at each age corresponds to the ex-ante distribution that any (unborn) individual faces. before learning his risk preferences and health outcomes.[34] Recall that our estimates imply that health and risk aversion are correlated, with more risk averse individuals being healthier on average. We assume that this correlation reflects differing stochastic processes for health conditional on an individual's innate risk aversion at age 25. To identify the stochastic health outcomes a 25-year old with a given risk aversion γ foresees at any given future age t, we isolate those individuals in our simulation sample (comprised of the 100 pseudo-individual versions of each person) of age t whose risk-aversion γ_t falls

[34] Recall that the age distribution in our sample is close to uniform, as it should be in a steady state population.

Average Costs at Various Ages Conditional on Age 25 Risk Aversion			
γ	30-35	45-50	55-60
0.0002	5,586	7,196	10,857
0.0003	4,212	6,390	10,319
0.0004	3,100	5,687	9,767
0.0005	2,328	4,911	9,271
0.0006	1,775	4,373	8,813

Table 7: Average costs as a function of age 25 risk preferences. Following the choice model estimates, costs are negatively related to risk aversion conditional on age.

into a band around the level expected based on our estimates of equation (5), for individuals with risk aversion level γ at age 25.[35] Table 7 shows for various risk aversion levels γ the average costs of the individuals selected in this manner at ages 25-30, 45-50, and 60-65. The pattern of costs reflects the positive correlation between health status and risk aversion, as well as the attenuation of this positive relationship with increases in age.

To measure the welfare difference for an individual with age-25 risk aversion level γ between any two regimes x and x' (under a given equilibrium concept e), we define the fixed yearly payment $y_{x,x'}(\gamma)$ added to income in regime x that makes the individual have the same expected utility starting at age 25 under regime x and as under regime x':

$$\sum_{t=25}^{65} \delta^t E[-e^{-\gamma\{I_t - CE_x(\lambda_t,\gamma)+y_{x,x'}(\gamma)\}}] = \sum_{t=25}^{65} \delta^t E[-e^{-\gamma\{I_t - CE_{x'}(\lambda_t,\gamma)\}}],$$

or

$$y_{x,x'}(\gamma) = -\frac{1}{\gamma}\ln\left(\frac{\sum_{t=25}^{65} \delta^t E[-e^{-\gamma\{I_t - CE_{x'}(\lambda_t,\gamma)\}}]}{\sum_{t=25}^{65} \delta^t E[-e^{-\gamma\{I_t - CE_x(\lambda_t,\gamma)\}}]}\right). \tag{7}$$

For a given discount factor $\delta \leq 1$ and regime x, we calculate $\sum_t \delta^t E[-e^{-\gamma\{I_t - CE_x(acg_t,\gamma)\}}]$ as follows: first, we first generate the value of $e^{-\gamma\{I_t - CE_x(\lambda_t,\gamma)\}}$ that each individual of age t in the band associated with γ would have if he chose between the 60 and 90 policies facing the equilibrium prices in regime x and having risk aversion parameter γ.[36] The income level I_t is either held fixed (in which case, with CARA preferences, its level doesn't matter) or comes from the regression in equation (6) and is estimated separately for managers and non-managers.[37] We then derive $E_{x_t}[-e^{-\gamma\{I_t - CE_x(\lambda_t,\gamma)\}}]$ by

[35] We use a band radius of 0.00005.

[36] Thus, we evaluate the welfare of an individual who at age 25 does not foresee his risk aversion changing.

[37] For managers, the mean income level I_t starts near income tier 1 at age 25 ($0-$40,000) and is near tier 4 at age 65 ($124,000-$176,000). Maximum income for managers occurs at age 66. For non-managers, mean income also starts near income tier 1 at age 25 and is halfway between tiers 2 ($40,000-$80,000) and 3 ($80,000-$124,000) at age 65. Income peaks at age 56 for non-managers, with an average near income tier 3. See the discussion of the estimates of equation (6) in Appendix C for more details.

Welfare Loss from Health Status–based Pricing in RE/sp-NE ($/year)			
γ	$y_{\lambda 4,no-pre}(\gamma)$ Fixed Income	$y_{\lambda 4,no-pre}(\gamma)$ Non-Manager Income path	$y_{\lambda 4,no-pre}(\gamma)$ Manager Income Path
0.0002	2,220	1,499	-384
0.0003	2,693	1,688	-613
0.0004	3,082	1,821	-886
0.0005	3,399	1,764	-973
0.0006	3,626	2,115	-891

Table 8: Long-run welfare comparison between the two pricing regulations of (i) pricing based on health status quartiles (x = "$\lambda 4$") and (ii) pure community rating / no pre-existing conditions (x' = "$no-pre$"). The table presents the values for $y_{x,x'}(\gamma)$, the annual payment required under regime x to make consumers indifferent between x and x'. The results presented are based on the RE and sp-NE equilibria outcomes presented in Table 5. We present results for the differing cases of (i) "fixed income" (ii) "income path" for non-managers and (iii) "income path" for managers. The assumed discount rate is $\delta = 0.975$.

calculating the sample mean of those values for age t individuals in the γ band. We then discount and add these values over t to get $\sum_t \delta^t E_{x_t}[-e^{-\gamma\{I_t - CE_x(\lambda_t,\gamma)\}}]$. We proceed similarly for regime x'.

We first compare two regimes: ACG-quartile pricing and pure community rating. The latter eliminates reclassification risk but exacerbates adverse selection. Health status-based pricing also involves some inter-temporal redistribution, as the young tend to face lower premiums. To the extent that this regime smooths consumption over time when borrowing is not possible (given the fact that income generally rises with age), this creates some welfare gain as well if agents cannot otherwise borrow to smooth their consumption over time.

Table 8 shows the values of $y_{x,x'}(\gamma)$ under the RE (and sp-NE) equilibrium notion comparing pricing based on ACG-quartiles (x = "$\lambda 4$") and no pre-existing conditions (x' = "$no - pre$"). We take $\delta = 0.975$. Since we do not know the extent to which agents are able to borrow to smooth their consumption, we compute welfare both assuming that income is fixed over time (perfect smoothing) and assuming they cannot borrow at all.[38] In the latter case we provide a calculation separately for managers and non-managers, whose expected incomes differ at each age.

With a fixed income, the welfare gains from eliminating reclassification when pricing pre-existing conditions is prohibited greatly exceed any losses this rule introduces due to adverse selection. The loss from health-based pricing on quartiles ranges from $2,220 to $3,626 per year depending on risk aversion level. Losses are larger for those with greater risk aversion. The annual loss with health status quartile pricing at a risk aversion level of 0.0004, approximately the mean in our sample, is $3,082, which is about 47% of the size of the $6,559 annual average total expenses in the population (see Table

[38] Our calculations do not, however, consider any gains from self-insurance through precautionary saving.

4 in Section 3). As a direct comparison point, we can compare this to the welfare implications of just adverse selection: with fixed income and risk aversion 0.0004 a consumer would be willing to pay \$619 per year to face a regime in which everyone receives the 90 policy at price $P_{90} = \underline{AC}_{90}$ rather than the community rating regime in which pre-existing conditions cannot be priced and everyone ends up buying the 60 policy at price $P_{60} = \overline{AC}_{60}$. Thus, the welfare loss from reclassification risk induced by this pricing regulation change is at least 5 times as large as the welfare loss from adverse selection under pure community rating.

When individuals cannot borrow, health-based pricing confers an additional benefit by moving consumption forward in life. For non-managers the losses from health-based pricing now range from \$1,499 to \$2,115 per year. For managers, however, whose income is higher and rises more steeply with age (see footnote 37), and therefore benefit more from moving consumption forward in time, health-based pricing is actually preferred to community rating. For this group, the benefits of smoothing income over time outweigh the costs of reclassification risk.

We revisit Table 6 to examine the welfare implications of varying the extent to which insurers can price health status information. Columns 3-5 illustrate the impact of finer pricing on long-run welfare. With fixed income (column 3), and for non-managers' income paths (column 4), the welfare loss from increased reclassification risk swamps the welfare gain from less adverse selection: the welfare loss from pricing 20 health status categories is almost 3 times that from pricing on quartiles. For managers' income paths the effect is not monotone, because of the benefits of income smoothing, but fine enough pricing does lead to a welfare loss relative to community rating (e.g., with 50 health status groups). Overall, the results highlight the trade-off between adverse selection and reclassification risk, and suggest that reclassification risk is likely to be more important from a welfare perspective.[39,40] We now turn to several extensions that examine policy-relevant modifications to our framework, e.g., age-based pricing, insurer risk-adjustment transfers designed to mitigate adverse selection, and the consequences of an unenforceable mandate.

6 Extensions

6.1 Age-Based Pricing

Age-based pricing is one of the few exceptions to pure community rating allowed in the Affordable Care Act. The legislation stipulates that insurers offering plans in any state-based exchange can vary prices on the basis of age by up to a 3:1 ratio: that is, older cohorts cannot be charged more than

[39] In addition to considering the fixed income case here, in the next section we consider the same comparison between community rating and pricing based on health status when there is also age-based pricing which eliminates the inter-temporal consumption-shifting effect of health status-based pricing. When we do so, managers also prefer community rating.

[40] One caveat to these results is that they rely on our estimated risk preferences being apprpropriate for evaluating reclassification risk; see the discussion in Section 6.5.

three times the premiums of younger buyers. States may further restrict this ratio to be lower (but cannot raise it) making this an important available policy tool. In this section, we use our framework to study market equilibrium and consumer welfare when insurers are allowed to vary prices with age and comment on (i) whether the 3:1 ratio is expected to bind, (ii) whether age-based pricing reduces adverse selection, and (iii) how the presence of age-based pricing affects the welfare impact of allowing pricing of pre-existing conditions. For a further investigation of age-based pricing regulation see e.g. Ericson and Starc (2013).

For this analysis, we group consumers into five-year age bins and use that categorization as the basis for differential pricing by insurers. Ericson and Starc (2013) shows that, in practice, insurers in Massachusetts price to five-year age bands, though this may or may not be the same in other states. Table 9 describes each age bucket and the equilibrium results. The first column shows mean total medical expenses by age in our sample: those age 30-35 have a mean of \$3,357 while those age 60-65 have a mean of \$9,413. Thus, just comparing average costs, the 3:1 age ratio appears to be non-binding in our sample judging by either ex-post or projected expenses captured by our health status measure (our upcoming analysis with MEPS data shows a similar pattern). If we include those age 25-30, whose mean expense is \$2,756, the age restriction is binding: however, in both the Affordable Care Act and Massachusetts legislation individuals in that age group can buy into a special catastrophic insurance pool that is separately regulated. As many of the healthiest individuals sign up for the catastrophic plan in Massachusetts, the 3:1 age restriction might not be binding in practice.[41] However, since policies typically include a deductible (e.g., our 60 policy involves a \$3,000 deductible) the relation between expected costs and the actuarial cost to the insurer is non-linear. The column on the right, labelled "Premium," reflects the actuarial cost of insuring each age group in the 60 plan given our plan designs. The deductible impacts the younger group the most, thus making the 3:1 ratio binding for the 30-35 population.

The second question we address is whether age-based pricing ameliorates the extent of adverse selection. As we saw in the previous section, by allowing some health status based pricing, additional trade was generated for the healthiest quartile of the population. Age — as shown in column 1 — is a proxy for health type, and may also enable more trade in equilibrium.

The final columns in Table 9 present equilibrium results with age-based pricing. Surprisingly perhaps, allowing for age-based pricing does not prevent full unraveling. For each age group, the RE/sp-NE involves all-in-60. Age-based pricing undoes some of the transfers from the younger, healthier age groups to the older groups that occur in pure community rating. However, the distributions of health risk still have substantial enough tails even for the younger age group that full unraveling occurs in equilibrium.[42]

[41] In Massachusetts, the age rating restriction is more restrictive than in the ACA, and doesn't allow age-based pricing at higher than a 2:1 ratio.

[42] We note that these results are robust to medium-sized changes in σ_ε, even though this shock to preferences introduces a source of willingness to pay for coverage unrelated to risk type. As we increase standard deviation of this shock, equilibria by age and for the whole population still involve unravelling to all-in-60. A σ_ε over 2,000 is required for some sub-markets

Age-based Pricing Regulation: Costs and Equilibrium Results						
	Expenses		ACG		Sh60	
Ages	Mean	S.D.	Mean	S.D.	RE/sp-NE	Premium
All	5,582	6,495	1.08	1.42	100.0	4,051
25-30	2,756	4,657	0.56	0.79	100.0	1,784
30-35	3,357	5,189	0.71	1.23	100.0	2,215
35-40	3,762	5,029	0.78	1.05	100.0	2,540
40-45	4,560	5,642	0.87	1.04	100.0	3,237
45-50	5,778	6,575	1.06	1.32	100.0	4,098
50-55	6,722	6,831	1.32	1.82	100.0	5,038
55-60	8,394	7,134	1.53	1.58	100.0	6,304
60-65	9,413	7,268	1.77	1.73	100.0	7,259

Table 9: Average costs, health status, and Riley /sp-NE equilibrium results for age-based rating policy.

Finally, we consider the simultaneous pricing of pre-existing conditions as well as age. Insurers will be allowed to price age in the exchanges, so it is natural to consider the effect of allowing pricing based on health status when age is priced as well. The exercise is interesting for at least two reasons. First, pricing based on health status may have a different impact on equilibrium in a more homogenous population, grouped by age, than it has in the whole population. Second, when evaluating the welfare impact of health status-based pricing, age-based pricing may neutralize the benefits associated with consumption smoothing, by reducing the transfer from young to old that health status-based pricing otherwise induces.

Table 10 shows the equilibrium when insurers can separate each age group into health status quartiles. Unlike pure age-based pricing which involved full unraveling to all-in-60 for every age group, we now have a positive share in 90 for all of the healthiest quartiles except in the oldest cohort, as well as for the second quartile for the younger groups. The interaction of age and health status-based pricing thus reduces adverse selection. Table 11 shows the compensation required to make an individual indifferent between a regime with health status quartile pricing for each age group, and another in which all individuals in each age band receive the 60 policy at its average cost for their age band. Once age is priced, health status-based pricing, which appealed to consumers with steeply increasing income, is no longer preferred by those consumers. The benefit of health status-based pricing is the reduction in adverse selection, and the postponement of premiums until later in life. With age-based pricing, the latter benefit is eliminated. The cost associated with reclassification risk then dominates the benefits of reducing adverse selection.

to not fully unravel.

Joint Health Status Quartile and Age Pricing Regulation: Equilibrium Results											
	Q1 (Healthy)			Q2			Q3		Q4 (Sick)		Avg.
Ages	Sh90	P60	P90	Sh90	P60	P90	Sh90	P60	Sh90	P60	Sh90
25-30	63	126	616	25	375	1,935	0	930	0	5,520	22
30-35	63	156	676	42	337	1,597	0	1,411	0	6,855	26
35-40	52	189	966	50	608	2,028	0	1,867	0	7,246	25
40-45	38	299	1,489	0	1,257	-	0	3,180	0	8,141	10
45-50	63	492	1,592	18	1,574	4,044	0	3,891	0	10,138	20
50-55	27	946	2,936	0	2,304	-	0	5,847	0	10,858	7
55-60	33	1,477	3,617	0	5,159	-	0	6,733	0	11,702	8
60-65	0	2,200	8,700	0	5,824	-	0	7,666	0	13,321	0

Table 10: RE /sp-NE equilibrium results for pricing regulations that allows insurers to price based on health status quartiles, conditional on age.

Welfare Loss from Health Status-quartile pricing, per age group, in RE/sp-NE ($/year)			
	$y_{\lambda 4+age,age}(\gamma)$	$y_{\lambda 4+age,age}(\gamma)$	$y_{\lambda 4+age,age}(\gamma)$
γ	Fixed Income	Non-Manager Income path	Manager Income Path
0.0002	2,111	2,129	1,100
0.0003	2,911	2,028	920
0.0004	3,707	1,842	778
0.0005	4,510	1,646	1,353
0.0006	5,137	1,612	1,876

Table 11: Long-run welfare comparison between the two pricing regulations of (i) pricing based on health status quartiles by age (x = "$\lambda 4 + age$")and (ii) pricing based on just age (x' = "age"). The results presented are based on the RE/sp-NE equilibria outcomes for each of the two pricing regulations. As before, the assumed discount rate is $\delta = 0.975$.

6.2 Participation and Subsidies

While the individual mandate will be a component of all exchanges to be implemented under the Affordable Care Act, in reality the mandate is a tax that is paid to the IRS when someone who can afford insurance remains uninsured. It is plausible that certain individuals, especially healthy ones, will decide to pay the mandated penalty and opt out. In Massachusetts, where an individual mandate has been in place since 2006, the penalty has been 50% of the cost of the least generous (Bronze) plan available through the exchange (Commonwealth Connector).[43] This is on average larger than the initial penalty under the Affordable Care Act, which is the maximum of \$695 per household member (up to three) and 2.5% of household income. In Massachusetts, only 3% of the population failed to comply with the individual mandate in 2008, with many of those people receiving exemptions from the penalty due to low income [Kolstad and Kowalski (2012)].

We simulate the role of the mandate. To do so we allow individuals to opt out of the exchanges should their expected utility from being uninsured be higher than joining their favorite insurance plan in the market. Uninsured means that the consumer pays zero premium and pays for the total cost of their health expenses.[44]

We present two exercises. First, we ask how the market would work absent the mandate. To answer that question we find equilibria allowing individuals to opt out without penalty. Second, we introduce opting out penalties, and determine equilibrium participation at different penalty levels.

Recall that equilibria without age-based pricing unravelled to all-in-60. The column "Better-off In" in the "Community Rating" section of Table 12 shows the percentage of each age group (and of the population as a whole) that is better off insured at the equilibrium premium of \$4,068 than remaining uninsured. For example, 44.2% ($= 100 - 55.8$) of 25 to 30 year old individuals prefer to opt out as their expected utility from non-insurance is higher than being pooled with the whole population.

Naturally, those that prefer to opt out are younger, healthier and less risk averse. The expected costs of insuring consumers who prefer to decline coverage is \$3,107 versus \$5,107 for those that prefer to participate. The average risk aversion coefficient of those that prefer to participate is $4.26 * 10^{-4}$ versus $4.03 * 10^{-4}$ for those that prefer to decline coverage.

Allowing healthier individuals to opt out increases the cost of covering the remaining pool, which in turn draws more people out of the pool. The process stops at a premium of \$5,339 when no more individuals want to drop out. At that premium there are no profitable single-policy Nash deviations in 60 or 90 to draw buyers back in. Thus, a P_{60} = \$5,339 is a sp-NE of the exchange without a mandate.[45] The equilibrium without the mandate involves full unravelling to 60, with 74.3% of the

[43] In 2010, in the 02138 zip code in Cambridge, MA, this penalty would have been \$5,500 for a family with two 40 year old parents, \$3,300 for a couple with two 35 year olds, and \$1,434 for a single 31 year old [Kolstad and Kowalski (2012)].

[44] This approach may overstate particpation for two reasons. First, as we discuss in Section 6.5, our estimated risk preferences may overstate aversion to the large risks involved with having no insurance. Second, as noted by Mahoney (2012), the choice to have no insurance may not mean that the individual pays for al of their medical expenses. On the other hand, other behavioral factors (social norms, etc.) may push individuals toward particpation.

[45] We have not examined whether these outcomes are RE or mp-NE.

	Implications of Individual Mandate					
	Community Rating		Age-Based Pricing			
	Mandate:	No Mandate:	Mandate:		No Mandate:	
Ages	Better-Off In	Participation	Premium	Better-Off In	Premium	Participation
All	78.3%	74.3%	-	80.7%	-	77.0%
25-30	55.8%	50.6%	1,786	70.1%	2,732	63.1%
30-35	59.6%	54.1%	2,215	70.0%	3,409	62.5%
35-40	68.7%	62.2%	2,542	75.9%	3,476	70.8%
40-45	75.1%	70.9%	3,242	77.7%	4,233	74.5%
45-50	82.5%	79.3%	4,103	82.9%	4,976	80.6%
50-55	90.6%	87.2%	5,038	88.6%	5,714	86.9%
55-60	94.7%	92.5%	6,304	92.1%	6,927	89.9%
60-65	95.8%	93.9%	7,259	91.6%	7,959	90.2%

Table 12: Implications of the individual mandate for equilibrium prices and market participation.

population voluntarily covered. The column "No Mandate: Participation" under "Community Rating" shows participation by age in the non-mandate equilibrium.

We can also compute the welfare impact of removing the mandate. Those individuals that remain covered, 74.3% of the population, suffer a loss equal to the premium increase \$1,271 ($= 5,339 - 4,068$). Comparing the certainty equivalent of remaining uninsured versus participation in the exchange for the 25.7% of the population that opts out, we find that they are better off by \$1,972, on average. Thus, removing the mandate entails a welfare loss of \$434.3 [$= 0.743(1,271) - 0.257(1,972)$] per person.

On the right side of Table 12 we show the corresponding numbers for age-based pricing. As we saw in Section 6.1, all the equilibria under the mandate (with no opting out) for the different age groups involve unravelling to 60. At the equilibrium premium, reported in the "Mandate: Premium" column, only a proportion of the population would voluntarily participate in the exchange. Column "Mandate: Better-off In," shows that the share that prefers to participate is an increasing share in age. Older individuals are more likely to benefit from participation, but the differences across ages are less pronounced once age is priced.

For each age, as individuals opt out, the cost of coverage increases. The column "No Mandate: Premium" reports the equilibrium premia for each age group absent a mandate. It is substantially higher than under the mandate, especially so for younger cohorts for whom the mandate is binding for a larger proportion of individuals.

We now turn to the second question: how do penalties impact the equilibrium? We now allow individuals to opt out if their expected utility from opting out and paying the penalty is higher than the utility from buying insurance at the equilibrium premium. (A participation subsidy plays the same

Market Participation as Function of Penalty		
Penalty	Participation %	Premium
0	74.3%	5,339
1000	78.2%	5,096
2000	82.7%	4,974
3000	88.4%	4,848
4000	96.7%	4,712
Free ins.	97.7%	4,144

Table 13: Market participation as a function of the penalty assessed for nonparticipation.

Age-based Pricing Regulation: Market Participation								
	30-35		40-45		50-55		60-65	
Penalty	Partic.	Prem.	Partic.	Prem.	Partic.	Prem.	Partic.	Prem.
0	63.1%	3,409	74.5%	4,232	86.9%	5,714	90.2%	7,959
1000	71.1%	3,047	79.3%	4,013	89.5%	5,570	91.8%	7,840
2000	87.6%	2,517	85.5%	3,755	92.5%	5,410	93.5%	7,717
3000	99.5%	2,225	94.9%	3,406	95.7%	5,252	95.3%	7,592

Table 14: Impact of individual mandate on participation and premiums by age group under age-based pricing.

role as the penalty.)

Table 13 shows participation for different penalty levels with community rating, and the corresponding equilibrium premium. While the penalty increases participation, it requires a very steep penalty (high subsidy) to even approach full participation. Good health draws are hard to persuade to remain in the pool. The last row presents participation under free insurance. Participation is not full, due to the preference shock.

Naturally, under age-based pricing the required penalty to achieve a specific participation level may be smaller. First, participation under age-based pricing is higher, but also healthier individuals who are on average younger might be easier to persuade to remain in a younger and healthier pool. Table 14 shows that a given subsidy/penalty is more effective in affecting participation when there is age-based pricing.

6.3 Risk Adjustment

In practice, the implementation of the ACA will involve risk adjustment transfers whose aim is to ameliorate adverse selection. States are free to propose their own risk adjustment rules, provided they receive Federal approval, or can instead default into the risk adjustment formula proposed by the

Federal government [see, e.g., Dept. of Health and Human Services (2012a) or Dept. of Health and Human Services (2012b)]. In this section, we use our model to evaluate the impact of the Federal formula. While in practice risk adjustment can lead to a number of problems, such as insurers up-coding enrollees to qualify for larger transfers, we will abstract from such issues and assume that the government can perfectly observe the health status of each enrollee.

It is tempting to think that risk adjustment can solve the adverse selection problem entirely, by simply providing a transfer to each firm that gives that firm an expected cost from each enrollee equal to the average cost if there was no selection, thereby "eliminating the impact of selection on cost." Unfortunately, doing so can result in the government running a deficit. As a result, the formula proposed by the HHS is designed to always break even. It provides a transfer payment per member to each plan i equal to

$$T_i = \left\{ \left(\frac{R_i}{\sum_i s_i R_i} \right) - \left(\frac{AV_i}{\sum_i s_i AV_i} \right) \right\} \overline{P}, \tag{8}$$

where R_i is plan i's "risk score" (equal to plan i's average cost divided by the average cost of all plans in the market), AV_i is plan i's actuarial value (i.e., 60 or 90 in our model), s_i is plan i's market share, and $\overline{P}$ is the average premium in the market. (Note that $\sum_i T_i = 0$, so the transfers are balanced.) These transfers alter insurers' average costs, which are now $AC_{90} - T_{90}$ and $AC_{60} - T_{60}$ in the 90 and 60 policy, respectively.

We will examine Riley equilibria. Since in equilibrium all policies break even and the transfers are balanced, the market average premium must equal the market average cost:[46]

$$\overline{P} = \overline{AC}(\Delta P) \equiv s_{90}(\Delta P) AC_{90}(\Delta P) + s_{60}(\Delta P) AC_{60}(\Delta P).$$

Plan i's risk score is $R_i = AC_{90}(\Delta P)/\overline{AC}(\Delta P)$.

Substituting into (8), we get

$$\begin{aligned} T_{90}(\Delta P) &= \left\{ \left(\frac{AC_{90}(\Delta P)}{\overline{AC}(\Delta P)} \right) - \left(\frac{0.9}{\overline{AV}(\Delta P)} \right) \right\} \overline{AC}(\Delta P) \\ &= AC_{90}(\Delta P) - \overline{AC}(\Delta P) \left(\frac{0.9}{\overline{AV}(\Delta P)} \right) \end{aligned}$$

where

$$\overline{AV}(\Delta P) \equiv s_{90}(\Delta P)(0.9) + s_{60}(\Delta P)(0.6).$$

[46] Formally, in equilibrium each policy will break even given its post-transfer average cost. Thus, recalling that T_i is a per member transfer, we have

$$P_{90} = AC_{90}(\Delta P) - T_{90}(\Delta P)$$

and

$$P_{60} = AC_{60}(\Delta P) + T_{90}(\Delta P) \left(\frac{s_{90}(\Delta P)}{s_{60}(\Delta P)} \right).$$

The market average premium is therefore

$$\overline{P} = s_{90}(\Delta P) P_{90} + s_{60}(\Delta P) P_{60} = \overline{AC}(\Delta P).$$

		Welfare Benefit of Risk-Adjusted Transfers: RE/sp-NE ($/year)	
γ	$y_{no-pre,risk-adj}(\gamma)$ Fixed Income	$y_{no-pre,risk-adj}(\gamma)$ Non-Manager Income path	$y_{no-pre,risk-adj}(\gamma)$ Manager Income Path
0.0001	316	261	106
0.0002	327	202	27
0.0003	336	139	18
0.0004	349	84	0
0.0005	368	36	38
0.0006	386	23	72

Table 15: Long-run welfare implications of insurer risk adjustment regulation (transfers based on risk mixture of the population enrolled).

Observe that the transfers depend on the market prices (through ΔP), while the market prices depend on the transfer rule. Thus, the equilibrium prices are determined as a fixed point. Specifically, the prices will be

$$\begin{aligned} P_{90}(\Delta P) &= AC_{90}(\Delta P) - T_{90}(\Delta P) \\ &= \overline{AC}(\Delta P)\left(\frac{0.9}{\overline{AV}(\Delta P)}\right) \end{aligned}$$

and

$$P_{60}(\Delta P) = \overline{AC}(\Delta P)\left(\frac{0.6}{\overline{AV}(\Delta P)}\right).$$

This leads to a fixed point condition for ΔP:

$$\Delta P = \overline{AC}(\Delta P)\left(\frac{0.3}{\overline{AV}(\Delta P)}\right). \tag{9}$$

Applying formula (9) to our data, we find that when health status-based pricing is not allowed the equilibrium with risk adjustment has prices $P_{90} = 6,189$ and $P_{60} = 4,139$, and the 90 policy capturing a 49% market share for the whole population.

To study the welfare implications we compare the long-run implications of equilibrium outcomes with and without insurer risk adjustment, for the case of pure community rating. Table 15 shows the yearly amount $y_{no-pre,risk-adj}$ an individual would be willing to pay to implement insurer risk adjustment relative to the case of pure community rating with no insurer transfers. The risk adjustment outcome is preferred, reflecting the reduction in adverse selection compared to the case with no insurer transfers.

6.4 Rebalancing of the Population

The analysis to this point has relied on health choice and utilization data from a large firm with approximately 10,000 employees and 20,000 covered lives. While these data have a lot of depth on dimensions that are essential to model health risk and risk preferences, they represent a specific population working for a specific large employer. Our results thus represent the case of exchange design as if this population were the population of interest. In this section we extend the analysis by applying our framework to a more externally relevant sample from the Medical Expenditures Panel Survey (MEPS), a survey that was specifically created to study medical care decisions for a nationally representative population. This analysis serves two purposes: (i) to study whether our broad conclusions are robust to changes in the population composition and (ii) to apply our analysis to a nationally representative sample.

We study individuals in the MEPS data from 2004 to 2008. The data are structured as overlapping two-year panels where each individual is in the panel for two consecutive years, and a new panel of individuals enters in each year. Table D.0 in Appendix D shows the number of individuals present in each year, which ranges from 33,066 to 34,403 over the five-year span. Note that because of the panel structure, an individual "counts" twice, once in each year they appear.

We base our analysis on two distinct samples that may be of interest to an exchange regulator, described in Table D.1 in Appendix D along with statistics for the full MEPS sample. Column 1 contains the summary statistics for the entire sample, with no sample cuts. Column 2 contains summary statistics for the sample of individuals between ages 25-65. In this table and in the analysis we split individuals within the same family into distinct individuals and run our simulations as if the entire market is an individual market.[47] Column 3 describes the sample of individuals between 25-65 who are uninsured, unemployed, or work for an employer that does not offer coverage (implying that if the individual has coverage, it is from the individual market). We run our exchange market analysis for the samples in Column 2 and Column 3. Column 2 is of interest if all individuals 25-65 participate in an exchange (e.g., the exchange insurers everybody). This sample is about half of the overall MEPS sample. Column 3 is of interest because it covers the population that will enter exchanges immediately when they are set up (uninsured and insured on individual market). This latter sample is about 15% of the overall MEPS sample, similar in age, lower in income, and more likely to come from the South. We note that this table presents the data "as is" while we use MEPS sample weights (which they use to make their own data representative) in our final analysis.

Table D.2 in Appendix D describes the insurance coverage for each of the three samples described in Table D.1. For individuals in Column 2 (age 25-65), 64% have some form of private insurance, 12% are on Medicaid, and 22% are uninsured. 76% of families have employer insurance offered at some point, while 62% always have an offer of employer insurance. Column 3, by design, has 83% of people

[47] Once the structure of family based premium setting is clear, we could run an additional analysis taking that into account, to the extent that it differs from aggregating up individual premiums into family premiums, which our current approach is essentially equivalent to.

uninsured and 17% of people insured on the individual market (with no employer insurance offer).

Our analysis matches individuals in the employer data used throughout our analysis to the two MEPS populations of interest (Columns 2 and 3) and creates two new simulation samples with demographic weights similar to the MEPS samples, but with detailed health and risk preference data from our estimates.[48] We match individuals in our data to those in the MEPS data based on three demographics: age, income, and gender. To do this, we probabilistically model cells of age, gender, and income in the MEPS samples, and then draw randomly from individuals in those bins in our data with weights proportional to the MEPS cell weights. We note that, before we construct the MEPS cell weights, we incorporate the sample weights in the MEPS data, which are intended to correct for sampling and response issues. Table D.3 in Appendix D describes the age, income, and gender cell multivariate cell weights for each of the two MEPS samples. We model the multivariate distribution of age and income fully non-parametrically, and assume that the probability of gender conditional on income is the same for all ages. We note that in this analysis, we do not match our sample to MEPS using health expenditure data (conditional on the other demographics) since our sample has more detailed medical information on consumers. However, the analysis below and the tables in Appendix D show that average costs conditional on demographics bins are similar in our data and in the MEPS data. Table D.4 provides more detail on the health risk for both MEPS samples.

Given this population weighting procedure, we study market equilibrium as in Section 4 for each weighted sample and each equilibrium concept. We study the pricing regimes of (i) pure community rating and (ii) health status quartile pricing, to compare to the results from our data. Table 16 presents the equilibrium simulation results for the sample re-weighted for all MEPS individuals from 25-65, while Table 17 presents the results for the sample re-weighted according to the MEPS uninsured / individual market population. These can be directly compared to Table 5. The comparison yields several important insights. First, the equilibrium premia and market shares are quite similar in the MEPS re-weighted samples to those from our main sample. For all samples the market fully unravels to all-in-60 for the case of pure community rating. For all samples, the healthiest quartile has substantial market share in both 60 and 90 under health status-based pricing: in our main analysis 64.8% in quartile 1 choose 60, for the re-weighted full MEPS analysis this value is 60.2%, while for the re-weighted uninsured / individual market MEPS sample it is 57.5%. Interestingly, there is also 30.4% of consumers enrolled in 90 for quartile 2 of the uninsured re-weighted sample, though for the other two samples this share is 0. This suggests that if the exchanges are comprised of only uninsured individuals and those that would have been on the individual market, there will be higher insurance rates for the within-exchange population. For the two sickest quartiles, everyone enrolls in 60 in equilibrium. Finally, and importantly, we note that the population expense levels are very similar between our main sample and the full re-weighted MEPS sample: if all enroll in 60, the average costs in the former are \$3,852 while in the latter they are \$4,051. For the uninsured re-weighted sample this value is \$3,901,

[48] We bring in the cost data from our data set because it is more detailed on the health risk dimension and our setting provides more precise plan characterizations, with which it is possible to estimate risk preferences.

MEPS Weighted: Equilibria without Pre-existing Conditions						
Equilirium Type	P_{60}	Sh_{60}	AC_{60}	P_{90}	Sh_{90}	AC_{90}
Riley / sp-NE	3,852	100.0	4,051	–	0	–
Multiple-policy Nash			Does not exist			
Wilson	3,774	78.5	2,112	6,674	21.5	12,760

MEPS Weighted: Equilibria with Health Status-based Pricing (Quartiles)							
Market	Equilibrium Type	P_{60}	Sh_{60}	AC_{60}	P_{90}	Sh_{90}	AC_{90}
Quartile 1	RE/sp-NE/mp-NE	321	60.2	321	1,521	39.8	1,521
Quartile 2	RE/sp-NE	1,445	100.0	1,445	-	0	-
Quartile 3	RE/sp-NE	4,239	100.0	4,239	-	0	-
Quartile 4	RE/sp-NE	9,347	100.0	9,347	-	0	-

Table 16: This table presents the analogous table to Table 5 on equilibrium outcomes, applied to the sample reweighted by characteristics of the MEPS full population, as described in the text. The top presents the equilibrium results for the case of pure community rating (no pricing of pre-existing conditions) and the bottom for the case where insurers can price based on health status quartiles.

but the distribution is more heavily skewed, with larger masses of both very healthy and very sick consumers. Overall, the analysis of MEPS data in this section suggests that, at a first pass, our main results are robust to different weighting of demographics to reflect a more nationally representative sample.

6.5 Lower Aversion toward Reclassification Risk

The welfare calculations above used the estimates of risk aversion from Section 3. Those estimates, however, were derived from situations in which individuals faced risks that were significantly smaller than the risks arising from reclassification, especially in the limit case where health status can be fully priced by insurers. Families choosing between the firm's PPO policies faced a maximum out-of-pocket cost of between $6,000 and $10,000, depending on their income, while reclassification risk can measure in the tens or even hundreds of thousands of dollars, depending on how much information insurers can use in pricing. The range of CARA risk aversion levels we estimated can imply implausible levels of risk aversion as the risks consumers consider are scaled up to higher monetary amounts [see, e.g., Rabin (2000)].

To address this issue we compute long-run welfare under a range of CRRA coefficients. The mac-

MEPS Unins. Weighted: Equilibria without Pre-existing Conditions						
Equilirium Type	P_{60}	Sh_{60}	AC_{60}	P_{90}	Sh_{90}	AC_{90}
Riley / sp-NE	3,901	100.0	3,901	–	0	–
Multiple-policy Nash			Does not exist			
Wilson	3,834	81.2	2,269	6,834	18.8	13,602

MEPS Unins. Weighted: Equilibria with Health Status-based Pricing (Quartiles)							
Market	Equilibrium Type	P_{60}	Sh_{60}	AC_{60}	P_{90}	Sh_{90}	AC_{90}
Quartile 1	RE/sp-NE/mp-NE	311	57.5	311	1,476	42.5	1,476
Quartile 2	RE/sp-NE	1,128	69.6	1,128	3,228	30.4	3,228
Quartile 3	RE/sp-NE	4,121	100.0	4,121	-	0	-
Quartile 4	RE/sp-NE	9,751	100.0	9,751	-	0	-

Table 17: This table presents the analogous table to Table 5 on equilibrium outcomes, applied to the sample reweighted by characteristics of the uninsured / individual coverage MEPS, described in the text. The top presents the equilibrium results for the case of pure community rating (no pricing of pre-existing conditions) and the bottom for the case where insurers can price based on health status quartiles.

Distributions of Certainty Equivalents of One-Year Risks				
	Community Rating		Health Status Quartile Pricing	
Age	Mean	S.D.	Mean	S.D.
30	5,840	1,037	3,455	3,581
45	6,622	1,067	5,659	4,034
60	7,580	1,030	9,389	4,111

Table 18: Distributions of certainty equivalents for one-year risk at different ages, from ex ante perspective.

roeconomics literature has considered a range of CRRA coefficients from 1 up to 10. Many authors concur that 1 to 4 are appropriate for the risks addressed in the finance and macro literature [see e.g. Barsky et al. (1997)]. However, the reclassification risk we model is substantially smaller than the risks considered in that literature. Table 18, for example, shows the mean and standard deviation of the distribution of certainty equivalents of the one-year risk faced under community rating and health status quartile pricing at the ages of 30, 45, and 60 for an individual whose CARA coefficient is 0.0004.[49] Under health status quartile pricing, the standard deviation of those certainty equivalents is roughly $4,000 at each age. That is $3,000 higher than with community rating.

Since we do not have a clear guidance on what risk parameter to use, we study the sensitivity of our findings to risk tolerance. As a reference, it is helpful to consider how an individual with CRRA preferences would respond to a bet with expected return of $1,000 (roughly the expected premium difference for a 40 year old), and a standard deviation of $4,000. An individual making $40,000 is indifferent between taking the bet and not taking it if their CRRA coefficient of risk aversion is 5. For $80,000 income, a coefficient of 10.1 is needed, and 15.2 if income is $120,000.

We can ask a similar question comparing the community rating and health status quartile regimes: at what level of CRRA would an individual be indifferent between them? For a fixed $75,000 income this would be a CRRA coefficient of 6.5, for an individual facing the non-manager income path it would be a coefficient of 9.5, and for someone with a manager's income path even a coefficient of 15 would lead them to prefer the health-status based pricing.[50]

6.6 Wilson Equilibria and Welfare

For robustness we also consider Wilson equilibria (WE). Wilson equilibria, as in the Rothschild-Stiglitz framework, may but need not coincide with the Riley and Nash equilibria. Indeed, we actually know that since the RE/sp-NE premiums we found do not survive a multi-policy Nash deviation, they are not a WE (since double deviations are unaffected by existing policies being dropped).

Wilson policies break even in total, but they do so allowing the 60 policy to cross-subsidize the 90 policy. We rely on the following result to identify WE in our data:

Proposition 3. *Let* $(\underline{P}_{90}^{BE}, \underline{P}_{60}^{BE})$ *be the break-even price configuration associated with* $\underline{\Delta P}^{BE}$*, and let* $\Delta P^w = Arg\max_{\Delta P \in [\underline{\theta}, \underline{\Delta P}^{BE}]} \Pi(\underline{P}_{60}^{BE} + \Delta P, \underline{P}_{60}^{BE})$*. If* $\Delta AC(\underline{\theta}) > \underline{\theta}$*, then the break-even price configuration* (P_{90}^w, P_{60}^w) *associated with price difference* ΔP^w *is a Wilson equilibrium.*

Thus, when $\Delta AC(\underline{\theta}) > \underline{\theta}$ (which is the case in our data), the price difference that maximizes the

[49] More precisely, for each age $t = 30$, 45, and 60, these are the distributions of the certainty equivalents of the one-year risk $CE_x(\lambda_t, \gamma)$ among those individuals in the year t band for age-25 risk aversion level $\gamma = 0.0004$. (These are the distributions we used in the welfare calculations in Section 5.)Thus, they reflect the uncertainty individuals face from an ex ante perspective.

[50] We do this calculation by using a 0.0004 CARA coefficient to derive the within-year certainty equivalents, and then apply the CRRA coefficient to the ex ante long-run utility calculation using those certainty equivalents.

Wilson Equilibria: Community Rating and Health Status-based Pricing (Quartiles)						
Market	P_{60}	Sh_{60}	AC_{60}	P_{90}	Sh_{90}	AC_{90}
Full Population	4,006	83.7	2,477	7,105	16.3	14,961
Quartile 1	302	60.2	290	1,502	39.8	1,519
Quartile 2	1,307	64.7	1,155	3,307	35.3	3,586
Quartile 3	4,443	70.0	3,337	7,193	30.0	9,648
Quartile 4	9,704	73.6	7,259	13,204	26.4	20,007

Table 19: Equilibrium results for Wilson solution concept for (i) pure community rating (no pre-existing conditions) and (ii) health status-based pricing with quartiles.

Welfare Loss from Health-Status-based Pricing (Quartiles) in Wilson Equilibrium ($/year)			
γ	$y_{\lambda 4,no-pre}(\gamma)$ Fixed Income	$y_{\lambda 4,no-pre}(\gamma)$ Non-Manager Income path	$y_{\lambda 4,no-pre}(\gamma)$ Manager Income Path
0.0002	2,101	1,390	-468
0.0003	2,577	1,592	-682
0.0004	2,964	1,711	-950
0.0005	3,277	1,628	-1,076
0.0006	3,506	1,923	-1,050

Table 20: Long-run welfare based on the Wilson Equilibrium results. Compares the two pricing regulations of (i) pricing based on health status quartiles (x = "$\lambda 4$") and (ii) pure community rating / no pre-existing conditions (x' = "$no-pre$").

profit from a multi-policy deviation from $(\underline{P}_{90}^{BE}, \underline{P}_{60}^{BE})$, the break-even price configuration associated with $\underline{\Delta P}^{BE}$, is a WE.[51]

Table 19 shows the equilibria with community rating and with health status quartile pricing. The cross-subsidization can be seen by comparing the prices to the average costs for each policy. We see that in every population the WE has a positive share of consumers purchasing the 90 policy, in contrast to the RE/sp-NE of Section 4.

Table 20 shows welfare results for WE. Here the welfare gains from prohibiting pricing based on pre-existing conditions with a fixed income are even larger, as the adverse selection losses from prohibiting pricing based on health status are smaller under the Wilson concept than under the Riley concept.

[51] In the Appendix, we also show that this is the unique Wilson equilibrium when $\Delta P^w \in (\underline{\theta}, \underline{\Delta P}^{BE})$. We conjecture, but have not proven that the same is true if $\Delta P^w \in \{\underline{\theta}, \underline{\Delta P}^{BE}\}$. We show as well that ΔP^w is also the price difference that leads to the lowest break-even P_{60} among all price differences at or below $\underline{\Delta P}^{BE}$.

7 Conclusion

In this paper we develop a model to study equilibrium and welfare for a class of regulated health insurance markets known as exchanges. We build on the theory of insurer competition under asymmetric information to develop a series of results characterizing equilibria in a stylized insurance exchange market, motivated by the exchanges proposed in the Affordable Care Act. We study several equilibrium concepts, including Riley equilibrium, Nash equilibrium, and Wilson equilibrium and prove results establishing the existence and uniqueness of Riley equilibrium in our environment, and provide an algorithm for identifying equilibria in our data. We estimate consumer micro-foundations, including ex ante health risk and risk preferences, using detailed health insurance choice and medical claims data from a large firm and use these estimates to simulate equilibria and long-run welfare using our data for a variety of possible pricing regulations.

The results reveal full unraveling in equilibrium with pure community rating, leading to all consumers being enrolled in the least comprehensive insurance available. As the health status information insurers can use to price their policies becomes finer and finer, the market share of the more comprehensive policy increases (less adverse selection) but consumers face greater year-to-year premium reclassification risk. We find that if consumers can borrow freely or if pricing based on age is also allowed (eliminating any consumption smoothing benefit of health-based pricing), the welfare loss from reclassification risk far outweighs the welfare gain from reduced adverse selection. Finally, we study a range of other market policies, including age-based pricing regulation, the implications of an unenforceable mandate, and insurer risk-adjustment. We also perform an analysis that matches our data to nationally representative MEPS data to illustrate that (i) this population is not that different from our own and (ii) the results are similar when we weight our data to reflect nationally representative demographic characteristics.

There are a number of dimensions on which our stylized model could be extended to more closely model most exchange environments. In our setting, products are differentiated only on financial dimensions. While in some settings (e.g., the Netherlands and Germany) this is essentially true in reality, in the U.S. context exchanges include insurers that offer products that are differentiated in terms of medical care and the network of available physicians. Accounting for this fact could enrich our equilibrium predictions and understanding of long-run welfare. In addition, it would be interesting to model more subtle consumer micro-foundations such as inertia or decision-making in complex product environments. With specific observed market regulations to guide us, it would be interesting to analyze markets where consumers can cover families as well as themselves (e.g., what are the implications of bundling individuals together into a contract). While these advances are exciting opportunities for future work, they require further developments of the model we develop which are beyond the scope of this paper.

More broadly, while our analysis has largely focused on the pricing restrictions in the Affordable Care Act in 2010, we have constructed a flexible framework that can be used to study a variety of counterfactual regulations. Besides the pricing regulations we study, important potential regulations to investigate include Minimum Creditable Coverage (the least comprehensive financial plans that can be offered) and contract regulation in general, which is a defining feature of exchange design. It would be particularly interesting to study dynamic, or long-run, insurance contracts that are more complex than the year-to-year contracts that we study but have the potential to reduce reclassification risk. While these kinds of contracts have been discussed to some extent [Cochrane (1995); Herring and Pauly (2006)], there has been little to no empirical analysis of such contracts, in part because they are complicated to analyze, both for researchers and consumers, and are observed infrequently in reality.

References

[1] Barsky, R.B., F.T. Juster, M.S. Kimball, and M.D. Shapiro (1993), "Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study," *Quarterly Journal of Economics* 112(2): 537-79.

[2] Bhattacharya, J., A. Chandra, M. Chernew, D. Goldman, A. Jena, D. Lakdawalla, A. Malani, and T. Phillipson (2013), "Best of Both Worlds: Uniting Universal Coverage and Personal Choice in Health Care," American Enterprise Institute.

[3] Bundorf, K., J. Levin, and N. Mahoney (2012), "Pricing and Welfare in Health Plan Choice," *American Economic Review* 102(7): 3214-3248.

[4] Capretta, J. and T. Miller (2010), "How to Cover Pre-Existing Conditions," , *National Affairs*, 4: 110-126.

[5] Cardon, J. and I. Hendel (2001), "Asymmetric Information in Health Insurance: Evidence from the National Medical Expenditure Survey," *RAND Journal of Economics* 32: 408-27.

[6] Carlin, C. and R. Town (2009), "Adverse Selection, Welfare, and Optimal Pricing of Employer Sponsored Health Plans," University of Minnesota working paper.

[7] Cochrane, J. (1995), "Time-Consistent Health Insurance," *Journal of Political Economy* 103: 445-73.

[8] Consumers Union. (2009),"What Will an Actuairal Value Standard Mean for Consumers?"

[9] Cutler, D. and S. Reber (1998), "Paying for Health Insurance: The Tradeoff Between Competition and Adverse Selection," *Quarterly Journal of Economics* 113: 433-66.

[10] Department of Health and Human Services. (2012a), "Patient Protection and Affordable Care Act: HHS Notice of Benefit and Payment Parameters for 2014; Proposed Rule: Part II."

[11] Department of Health and Humans Services. (2012b), "Risk Adjustment Methodology Overview."

[12] Einav, L. and A. Finkelstein (2011), "Selection in Insurance Markets: Theory and Empirics in Pictures," *Journal of Economic Perspectives* 25(1): 115-38.

[13] Einav, L., A. Finkelstein, and J. Levin (2010a), "Beyond Testing: Empirical Models of Insurance Markets," *Annual Review of Economics* 2: 311-36.

[14] Einav, L., A. Finkelstein, and P. Schrimpf (2010b), "Optimal Mandates and the Welfare Cost of Asymmetric Information: Evidence from the U.K. Annuity Market," *Econometrica* 78: 1031-92.

[15] Einav, L., A. Finkelstein, and M. Cullen (2010c), "Estimating Welfare in Insurance Markets Using Variation in Pricing," *Quarterly Journal of Economics* 125(3): 877-921.

[16] Einav, L., A. Finkelstein, S. Ryan, P. Schrimpf, and M. Cullen (2013), "Selection on Moral Hazard in Health Insurance," *American Economic Review* 103(1): 178-219.

[17] Engers, M. and L. Fernandez (1987), "Market equilibrium with Hidden Knowledge and Self-selection," *Econometrica* 55: 425-39.

[18] Ericson, K. and A. Starc. (2013), "Pricing Regulation and Imperfect Competition on the Massachusetts Health Insurance Exchange," Wharton working paper.

[19] Fernandez, B., and A. Mach (2012), "Health Insurance Exchanges Under the Patient Protection and Affordable Care Act," Congressional Research Service.

[20] Finkelstein, A. and K. McGarry (2006), "Multiple Dimensions of Private Information: Evidence from the Long-Term Care Insurance Market," *American Economic Review* 96: 938-958.

[21] Finkelstein, A., K. McGarry, and A. Sufi, A. (2005), "Dynamic Inefficiencies in Insurance Markets: Evidence from Long-Term Care Insurance," MIT working paper.

[22] Handel, B. (2013), "Adverse Selection and Switching Costs in Health Insurance Markets: When Nudging Hurts," UC-Berkeley working paper.

[23] Hendel, I. and A. Lizzeri (2003), "The Role of Commitment in Dynamic Contracts: Evidence from Life Insurance," *Quarterly Journal of Economics* 118: 299-327.

[24] Herring, B. and M. Pauly (2006), "Incentive-Compatible Guaranteed Renewable Health Insurance Premiums," *Journal of Health Economics* 25: 395-417.

[25] Kaiser Family Foudnation. (2010), "Focus on Health Reform: Summary of New Health Reform Law," http://www.kff.org/healthreform/upload/8061.pdf.

[26] Koch, T. (2011), "One Pool to Insure Them All? Age, Risk, and the Price(s) of Medical Insurance," UC-Santa Barbara Working Paper.

[27] Kolstad, J. and A. Kowalski (2012), "Mandate-Based Health Reform and Evidence from the Labor Market: Evidence from the Massachusetts Reform," Wharton working paper.

[28] Lustig, J. (2010), "Measuring Welfare Losses from Adverse Selection and Imperfect Competition in Privatized Medicare," Boston University working paper.

[29] Mahoney, N. (2012), "Bankrupcy as Implicit Insurance," working paper.

[30] Miyazaki, H. (1977), "The Rat Race and Internal Labor Markets," *Bell Journal of Economics* 8: 394-418.

[31] Rabin, M. (2000), "Risk Aversion and Expected-Utility Theory: A Calibration Theorem," *Econometrica* 68(5): 1281-92.

[32] Riley, J. G. (1985), "Informational Equilibrium," *Econometrica* 47: 331-59.

[33] Rothschild, M. and J. E. Stiglitz (1976), "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," *Quarterly Journal of Economics* 90: 629-49.

[34] Train, K. (2009), *Discrete Choice Methods with Simulation, 2nd. Ed.*, Cambridge University Press.

[35] Wilson, C. (1977), "A Model of Insurance Markets with Imperfect Information," *Journal of Economic Theory* 16: 167-207.

A Appendix: Proofs

In what follows, we refer to the lowest offered prices (P_{90}, P_{60}) for the two policies as a "price configuration." An "outcome" refers to the market shares of the two policies and the lowest offered prices for any policies that are purchased by a strictly positive share of consumers.

A.1 Nash Equilibria

We first discuss Nash equilibria (NE), both with single-policy and multi-policy firms. Note that we can without loss of generality restrict attention to Nash equilibria in which more than one firm offers each of the prices in an equilibrium price configuration. To see this, suppose we have a Nash equilibrium, with equilibrium price configuration (P_{90}^*, P_{60}^*), in which this is not true. Introducing additional firms that offer these prices makes the profit earned by any active firm through a deviation identical to the profit earned by an entrant making the same deviation. Since no entrant wanted to deviate in the initial Nash equilibrium despite having zero profits (which is weakly less than the profit of any active firm in the equilibrium), no active firm will want to deviate at this new set of price offers. Moreover, no entrant will want to deviate either (an entrant's deviation profit hasn't changed), so the new set of price offers is a Nash equilibrium and results in an identical price configuration and outcome as in the initial equilibrium. Finally, note that we can therefore test whether a price configuration arises in a Nash equilibrium by examining solely the profitability of entrant deviations.

The discussion that follows establishes Proposition 1.

A.1.1 Single-policy firms

Lemma 1. *Suppose that (P_{90}^*, P_{60}^*) is a NE price configuration. Then both policies must break even: i.e.,* $\Pi_{90}(P_{90}^*, P_{60}^*) = \Pi_{60}(P_{90}^*, P_{60}^*) = 0$.

Proof. Suppose policy k is profitable. A new firm could enter and offer price $P_k - \varepsilon$, where $\varepsilon > 0$, for that contract and make a positive profit: This deviation attracts all consumers who are purchasing policy k, and earns a positive profit on them, and for small ε, attracts very few others. Taking $\varepsilon \to 0$ yields the result. $\square$

Lemma 2. *Among all price configurations (P_{90}, P_{60}) at which both policies break even and there are positive sales of the 60 policy ($\Delta P > \underline{\theta}$), only the one with the lowest sales of the 60 policy (i.e., having $\Delta P = \underline{\Delta P}^{BE}$) can be NE.*[52]

Proof. Suppose that at price configurations $P' = (P'_{90}, P'_{60})$ and $P'' = (P''_{90}, P''_{60})$ both policies break even, $\min\{\Delta P', \Delta P''\} > \underline{\theta}$, and there is a larger share for the 60 policy in P'' than in P'. Then $\Delta P' < \Delta P''$ and there are positive sales of the 90 policy at P'.[53] In addition, $P'_{60} = AC_{60}(\Delta P') < AC_{60}(\Delta P'') =$

[52] Such a policy configuration may have all consumers buying the 60 policy.

[53] Note that we may have $\Delta P'' = \overline{\theta}$.

P''_{60}. Starting at price configuration P'', consider a deviation in which an entrant offers price $\widehat{P}_{90} \equiv P''_{60} + \Delta P' < P''_{60} + \Delta P'' = P''_{90}$. Since $\widehat{P}_{90} - P''_{60} = \Delta P'$, after the deviation the share of the 90 policy is the same (positive) share as at P'. Moreover,

$$\begin{aligned}\widehat{P}_{90} - AC_{90}(\widehat{P}_{90} - P''_{60}) &= (P''_{60} + \Delta P') - AC_{90}(\Delta P') \\ &> (P'_{60} + \Delta P') - AC_{90}(\Delta P') \\ &= P'_{90} - AC_{90}(\Delta P') \\ &= 0.\end{aligned}$$

Thus, the deviation is profitable, implying that P'' is not a Nash equilibrium price configuration. □

Lemma 2 tells us that we can narrow down the possible equilibria to two: (i) everyone in the 90 policy with $P_{90} = \underline{AC}_{90}$ and $\Delta P = \underline{\theta}$, and (ii) $\Delta P = \underline{\Delta P}^{BE}$ (the lowest break-even $\Delta P \in (\underline{\theta}, \overline{\theta}]$, at which there are positive sales of the 60 policy). The former ("all-in-90") is an equilibrium only if there is no profitable deviation to a $P_{60} < \underline{AC}_{90} - \underline{\theta}$. Moreover, observe the following:

Lemma 3. *If there is a NE outcome in which all consumers buy the 90 policy, then this is the unique NE outcome.*

Proof. By Lemma 1, the equilibrium price of the 90 policy in this NE price configuration, say $P^* = (P^*_{90}, P^*_{60})$, has $P^*_{90} = \underline{AC}_{90}$. Suppose as well that there is a NE price configuration $P^{**} = (P^{**}_{90}, P^{**}_{60})$ in which policy 60 is purchased by a positive measure of consumers. By Lemma 1, we must therefore have $P^{**}_{60} > \underline{AC}_{60}$. If $\underline{AC}_{90} - \underline{\theta} \geq P^{**}_{60}$, then starting from the all-in-90 price configuration P^*, a deviation to P^{**}_{60} would attract all consumers to the 60 policy and be strictly profitable (since $P^{**}_{60} > \underline{AC}_{60}$), breaking the all-in-90 equilibrium. So, it must be that $\underline{AC}_{90} - \underline{\theta} < P^{**}_{60}$. But in that case, starting from the price configuration P^{**}, an entrant could earn a strictly positive profit by offering the 90 policy at price $\widehat{P}_{90} \equiv P^{**}_{60} + \underline{\theta}$. This deviation attracts all of the consumers, and earns a positive profit since

$$\widehat{P}_{90} - AC_{90}(\widehat{P}_{90} - P^{**}_{60}) = (P^{**}_{60} + \underline{\theta}) - \underline{AC}_{90} > 0.$$

Thus, P^{**} cannot be an equilibrium price configuration – a contradiction. □

These results suggest the following procedure: First, see if all consumers buying the 90 policy is a NE by checking if there is profitable deviation in P_{60}. If it is a NE, then it is the unique NE. If it is not a NE, then identify the price difference $\underline{\Delta P}^{BE}$ (the lowest break-even ΔP with positive sales of the 60 policy) and check if the price configuration associated with $\underline{\Delta P}^{BE}$ is a NE. In applying this procedure, two useful results are the following:

Lemma 4. *If $\Delta AC(\underline{\theta}) > \underline{\theta}$, then there cannot be a NE outcome in which all consumers buy the 90 policy.*

Proof. In an all-in-90 equilibrium, we must have $P_{90} = \underline{AC}_{90}$. A deviation in P_{60} to $\widehat{P}_{60} = \underline{AC}_{90} - \underline{\theta} - \varepsilon$ for sufficiently small $\varepsilon > 0$ attracts a positive measure of consumers with types close to $\underline{\theta}$ and earns strictly positive profits since $\widehat{P}_{60} - \underline{AC}_{60} = (\underline{AC}_{90} - \underline{\theta}) - \underline{AC}_{60} = \Delta AC(\underline{\theta}) - \underline{\theta} > 0$. □

Lemma 5. *Suppose that at price configuration (P_{90}, P_{60}) we have $\Pi_{90}(P_{90}, P_{60}) = \Pi_{60}(P_{90}, P_{60}) = 0$ and positive sales of the 60 policy. Then no single-policy deviation by an entrant in P_{60} is profitable.*

Proof. An entrant's deviation $\widehat{P}_{60}$ must be below P_{60} to attract any consumers. But any such deviation makes losses, since

$$\widehat{P}_{60} - AC_{60}(P_{90} - \widehat{P}_{60}) < P_{60} - AC_{60}(P_{90} - P_{60}) = 0.$$

□

Lemma 4 provides a simple sufficient condition for all-in-90 to fail to be an equilibrium, while Lemma 5 tells us that we can restrict attention to deviations in P_{90} when checking if $\underline{\Delta P}^{BE}$ arises in a NE.

A.1.2 Multi-policy firms (with multi-policy deviations)[54]

Clearly, Lemmas 1-5 still hold (note that all of the deviations used in the proofs of those results are still feasible). The new result is as follows:

Lemma 6. *If $P^* = (P^*_{90}, P^*_{60})$ is a NE price configuration with multi-policy firms, then*

(i) if all consumers are buying the 90 policy, P^* is a NE price configuration if there is no profitable single policy deviation by an entrant in P_{60};

Proposition 4.

(ii) *if some consumers are buying the 60 policy, P^* is a NE price configuration iff $\Pi(P^*_{90}, P^*_{60}) = 0 = \max_{\widehat{P}_{90} \leq P^*_{90}} \Pi(\widehat{P}_{90}, P^*_{60})$, that is, if there is no profitable multi-policy deviation by an entrant that reduces P_{90} and lowers P_{60} slightly to capture all consumers.*

Proof. For part (i), let (P^*_{90}, P^*_{60}) be the candidate equilibrium in which only the 90 policy is purchased (at price $P^*_{90} = \underline{AC}_{90}$). Consider any entrant multi-policy deviation $(\widehat{P}_{90}, \widehat{P}_{60}) \leq (P^*_{90}, P^*_{60})$. To be profitable, some consumers must buy policy 60 in the deviation, so $\widehat{P}_{60} < P^*_{60}$ and $\Delta\widehat{P} > \Delta P^*$. But the most profitable such deviation must have $\widehat{P}_{90}$ equal to or arbitrarily close to P^*_{90}. (Otherwise, both $\widehat{P}_{90}$ and $\widehat{P}_{60}$ could be raised by a small and equal amount.) But, since the reduction in P_{60} makes the 90 policy at price P^*_{90} unprofitable, this deviation is less profitable for the entrant than a single-policy deviation to $\widehat{P}_{60}$.

[54] Note that when firms can offer multiple policies but can deviate in only one policy at a time there is no change from the analysis of the previous subsection.

For part (ii), we already know that an entrant's single-policy deviation in P_{60} is unprofitable (Lemma 5). A single-policy deviation offering the 90 policy at price $\widehat{P}_{90} < P^*_{90}$, since it makes the 60 policy at price P^*_{60} earn strictly positive profits, is less profitable than the multi-policy deviation $(\widehat{P}_{90}, P^*_{60} - \varepsilon)$ for sufficiently small $\varepsilon > 0$, as this captures the entire market. As $\varepsilon \to 0$, this deviation has profits equal to $\max_{\widehat{P}_{90} \leq P^*_{90}} \Pi(\widehat{P}_{90}, P^*_{60})$. $\square$

Although it will not pay a role in our analysis, we note the following result:

Lemma 7. *If $\theta > C_{90}(\theta) - C_{60}(\theta)$ for all $\theta \in [\underline{\theta}, \overline{\theta}]$, then some consumers must be buying the 90 policy in any NE with multi-policy firms.*

Proof. Suppose all consumers were purchasing the 60 policy. Then, by Lemma 1, $P^*_{60} = \overline{AC}_{60}$ and $P^*_{90} = P^*_{60} + \overline{\theta}$. Now consider a deviation to $(P^*_{90} - \varepsilon, P^*_{60})$. We will show that for small $\varepsilon > 0$, aggregate profits are strictly positive. Aggregate profits equal

$$\psi(\varepsilon) \equiv \Pi(P^*_{90} - \varepsilon, P^*_{60}) = \int_{\overline{\theta}-\varepsilon}^{\overline{\theta}} [P^*_{90} - \varepsilon - C_{90}(\theta)] f(\theta) d\theta + \int_{\underline{\theta}}^{\overline{\theta}-\varepsilon} [P^*_{60} - C_{60}(\theta)] f(\theta) d\theta.$$

Now

$$\psi'(\varepsilon) = [P^*_{90} - \varepsilon - C_{90}(\overline{\theta} - \varepsilon)] f(\overline{\theta} - \varepsilon) - [P^*_{60} - C_{60}(\overline{\theta} - \varepsilon)] f(\overline{\theta} - \varepsilon) - [1 - F(\overline{\theta} - \varepsilon)],$$

so

$$\begin{aligned} \psi'(0) &= [P^*_{90} - C_{90}(\overline{\theta})] f(\overline{\theta}) - [P^*_{60} - C_{60}(\overline{\theta})] f(\overline{\theta}) \\ &= f(\overline{\theta}) \{\overline{\theta} - [C_{90}(\overline{\theta}) - C_{60}(\overline{\theta})]\} > 0. \end{aligned}$$

Since, by Lemma 1, $\psi(0) = \Pi(P^*_{90}, P^*_{60}) = 0$, this implies that for small $\varepsilon > 0$ aggregate profit is strictly positive. As a result, there is a $\delta > 0$ such that $(P^*_{90} - \varepsilon, P^*_{60} - \delta)$ is a profitable deviation. $\square$

The assumption that $\theta > C_{90}(\theta) - C_{60}(\theta)$ for all $\theta \in [\underline{\theta}, \overline{\theta}]$ is an implication of risk aversion; it says that all consumers prefer the greater coverage of the 90 policy if it is priced at fair odds (for that consumer). However, in our analysis the presence of a (behavioral) idiosyncratic preference shock for each policy could mean that consumers do not satisfy this condition.

A.2 Riley Equilibria

A.2.1 Safe price offers

We begin by considering which price offers are "safe" in the sense that they do not incur losses regardless of any additional offers being introduced.

Lemma 8. *Given price configuration (P_{90}, P_{60}), single-policy offer $P''_{60} < P_{60}$ is safe if and only if $\Pi_{60}(P_{90}, P''_{60}) \geq 0$.*

Proof. If $\Pi_{60}(P_{90}, P''_{60}) < 0$, then P'_{60} makes losses absent any reaction, and hence is not safe. So suppose that $\Pi_{60}(P_{90}, P''_{60}) \geq 0$. Any price offers $\widehat{P} = (\widehat{P}_{90}, \widehat{P}_{60})$ with a $\widehat{P}_{60} < P''_{60}$ gives the firm offering P''_{60} a profit of zero. Any price offers $\widehat{P}$ with $\widehat{P}_{90} \geq P_{90}$ and $\widehat{P}_{60} \geq P''_{60}$ cannot make the firm offering P''_{60} incur losses. Finally, any price offers $\widehat{P}$ with $\widehat{P}_{90} < P_{90}$ and $\widehat{P}_{60} \geq P''_{60}$ weakly lowers the sales of the firm offering P''_{60}. If that firm makes no sales at $(\widehat{P}_{90}, P''_{60})$, then its profit is zero. If it has positive sales at $(\widehat{P}_{90}, P''_{60})$, then it must also at (P_{90}, P''_{60}). This implies that $\Pi_{60}(\widehat{P}_{90}, P''_{60}) \geq 0$ since then $AC_{60}(\widehat{P}_{90} - P''_{60}) \leq AC_{60}(P_{90} - P''_{60}) \leq P''_{60}$. □

Definition 2. *The* ***lowest safe 60 price given*** P_{90} *is* $\underline{P}_{60}(P_{90}) \equiv \min\{P''_{60} : \Pi_{60}(P_{90}, P''_{60}) \geq 0\}$.

Remark 1. *Define the price* $\widetilde{P}_{60}(P_{90}) \equiv \{\ \widetilde{P}_{60} : \widetilde{P}_{60} = AC_{60}(P_{90} - \widetilde{P}_{60})\}$. *Note that this equality has a unique solution, which is continuous and weakly increasing in* P_{90} *and strictly increasing at any* P_{90} *at which* $P_{90} - \widetilde{P}_{60}(P_{90}) \in (\underline{\theta}, \overline{\theta})$ *(so that there are sales of both policies). Since* $\widetilde{P}_{60}(\underline{AC}_{60} + \underline{\theta}) = \underline{AC}_{60}$ *and* $\widetilde{P}_{60}(\overline{AC}_{60} + \overline{\theta}) = \overline{AC}_{60}$, *this occurs when* $P_{90} \in (\underline{AC}_{60} + \underline{\theta}, \overline{AC}_{60} + \overline{\theta})$. *See Figure 1. Moreover, one can see in the figure that* $AC_{60}(P_{90} - \widetilde{P}_{60}(P_{90}))$ *is strictly increasing for this range of* P_{90}, *which also means that* $P_{90} - \widetilde{P}_{60}(P_{90})$ *is strictly increasing, and these are weakly increasing at all* P_{90}. *Moreover, observe that the lowest safe 60 price given* P_{90} *is given by the continuous function:*

$$\underline{P}_{60}(P_{90}) = \left\{ \begin{array}{cc} P_{90} - \underline{\theta} & \text{if } P_{90} \leq \underline{AC}_{60} + \underline{\theta} \\ \widetilde{P}_{60}(P_{90}) & \text{if } P_{90} \in (\underline{AC}_{60} + \underline{\theta}, \overline{AC}_{60} + \overline{\theta}) \\ \overline{AC}_{60} & \text{if } P_{90} \geq \overline{AC}_{60} + \overline{\theta} \end{array} \right\}.$$

When $P_{90} \leq \underline{AC}_{60} + \underline{\theta}$, *all consumers buy the 90 policy at prices* $(P_{90}, \underline{P}_{60}(P_{90}))$; *when* $P_{90} \in (\underline{AC}_{60} + \underline{\theta}, \overline{AC}_{60} + \overline{\theta})$ *there are positive sales of both policies at prices* $(P_{90}, \underline{P}_{60}(P_{90}))$; *and when* $P_{90} \geq \overline{AC}_{60} + \overline{\theta}$ *all consumers buy the 60 policy at prices* $(P_{90}, \underline{P}_{60}(P_{90}))$.

Remark 2. *Observe that if a two-policy reaction* (P''_{90}, P''_{60}) *is safe and causes the profitable single-policy deviation* P'_{90} *to instead make losses, then the single-policy reaction* P''_{60} *is also safe and causes the single-policy deviation* P'_{90} *to make losses. To see why, note first that it cannot be that* $P''_{90} < P'_{90}$ *(otherwise the deviator's profit would not be strictly negative). The result is immediate if* $P''_{90} > P'_{90}$. *So suppose that* $P''_{90} = P'_{90}$. *Since the firms make losses on the 90 policy and the reaction is safe, we must have* $\Pi_{60}(P'_{90}, P''_{60}) > 0$. *But then Lemma 8 implies that the single-policy reaction* P''_{60} *is safe and clearly also causes the deviating firm to make losses. Hence, in looking at safe reactions to single-policy deviations in* P_{90}, *we can restrict attention to single-policy safe reactions in* P_{60}.

Lemma 9. *If at* $(P_{90}, \underline{P}_{60}(P_{90}))$ *we have positive sales of the 90 policy and* $\Pi_{90}(P_{90}, \underline{P}_{60}(P_{90})) \geq 0$, *then* $\Pi_{90}(P_{90}, P_{60}) \geq 0$ *at all* $P_{60} > \underline{P}_{60}(P_{90})$.

Proof. Since there are positive sales of the 90 policy, it follows that $P_{90} \geq AC_{90}(P_{90} - \underline{P}_{60}(P_{90})) \geq AC_{90}(P_{90} - P_{60})$ for any $P_{60} > \underline{P}_{60}(P_{90})$, where the second inequality follows from that fact that increases in P_{60} weakly lower AC_{90}. □

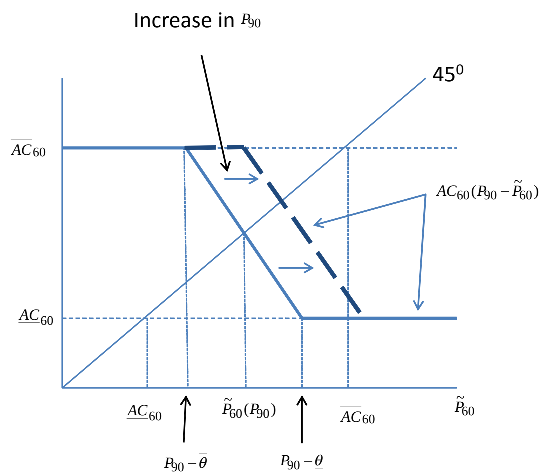

Figure A1: Graphical Description of the Lowest Safe P_{60} given P_{90}.

Remark 3. *In light of Remark 2, Lemma 9 implies that a profitable single-policy deviation to P'_{90} can be rendered unprofitable by a safe reaction if and only if it is rendered unprofitable by a single-policy reaction to $\underline{P}_{60}(P'_{90})$.*

A.2.2 Characterization of Riley Equilibria

Just as with Nash equilibria, we can restrict attention to Riley equilibria (RE) in which each price in the equilibrium price configuration is offered by more than one firm. We establish Proposition 2 through a series of lemmas.

Lemma 10. *If (P^*_{90}, P^*_{60}) is a RE, then $\Pi_{90}(P^*_{90}, P^*_{60}) = \Pi_{60}(P^*_{90}, P^*_{60}) = 0$.*

Proof. We first show that $\Pi_{60}(P^*_{90}, P^*_{60}) \leq 0$. Suppose otherwise, so that $\Pi_{60}(P^*_{90}, P^*_{60}) > 0$. Then for small $\varepsilon > 0$ we would have $\Pi_{60}(P^*_{90}, P^*_{60} - \varepsilon) > 0$. By Lemma 8, a single-policy deviation that offers $P^*_{60} - \varepsilon$ would then be safe, and there would therefore be no reaction that could render it unprofitable. But then (P^*_{90}, P^*_{60}) would not be a Riley equilibrium, a contradiction.

We next show that $\Pi_{90}(P^*_{90}, P^*_{60}) \leq 0$. The result is immediate if $\Delta P^* \equiv P^*_{90} - P^*_{60} = \overline{\theta}$ so that the 90 policy makes no sales at (P^*_{90}, P^*_{60}). So suppose that $\Delta P^* < \overline{\theta}$ (implying that the 90 policy has positive sales) and that contrary to the claim (P^*_{90}, P^*_{60}) is a Riley equilibrium with $\Pi_{90}(P^*_{90}, P^*_{60}) > 0$. If $\underline{P}_{60}(P^*_{90}) > P^*_{60}$, then a single-policy deviation to $P^*_{90} - \varepsilon$ for small enough $\varepsilon > 0$ would be a profitable Riley deviation as no safe reaction in P_{60} could render it unprofitable. So we must have $\underline{P}_{60}(P^*_{90}) \leq P^*_{60}$.

Now if $\underline{P}_{60}(P^*_{90}) < P^*_{60}$, then there can be no 60 sales at (P^*_{90}, P'_{60}) for any $P'_{60} \in [\underline{P}_{60}(P^*_{90}), P^*_{60})$, since otherwise a single-policy deviation to $P^{'}_{60} + \varepsilon$ for sufficiently small $\varepsilon > 0$ would be profitable and safe. Thus, $\underline{P}_{60}(P^*_{90}) \leq P^*_{60}$ implies that $\Pi_{90}(P^*_{90}, \underline{P}_{60}(P^*_{90})) = \Pi_{90}(P^*_{90}, P^*_{60}) > 0$. By continuity, we then have that $\Pi_{90}(P^*_{90} - \varepsilon, \underline{P}_{60}(P^*_{90} - \varepsilon)) > 0$ for small enough $\varepsilon > 0$, so a single-policy deviation to such a $P^*_{90} - \varepsilon$ cannot be rendered unprofitable by any safe reaction, yielding a contradiction.

Thus, we have $\Pi_{60}(P^*_{90}, P^*_{60}) \leq 0$ and $\Pi_{90}(P^*_{90}, P^*_{60}) \leq 0$. But if either is strictly negative, then some firm must be earning strictly negative profits, and would do better by dropping all of its policies. The result follows. □

Lemma 11. *There is a Riley equilibrium in which all consumers buy the 90 policy if and only if there is no P'_{60} such that $\Pi_{60}(\underline{AC}_{90}, P'_{60}) > 0$.*

Proof. By Lemma 10, $P^*_{90} = \underline{AC}_{90}$. Necessity follows because if there was a P'_{60} such that $\Pi_{60}(\underline{AC}_{90}, P'_{60}) > 0$, then a single-policy deviation to P'_{60} would be profitable and safe (Lemma 8), so could not be rendered unprofitable by any reaction. For sufficiency, suppose that there is no P'_{60} such that $\Pi_{60}(\underline{AC}_{90}, P'_{60}) > 0$ but that $P' = (P'_{90}, P'_{60})$ is a profitable Riley deviation. If the deviator makes no sales of the 90 policy, then its sales would be the same with a single-policy deviation to P'_{60}, so we would have $\Pi_{60}(\underline{AC}_{90}, P'_{60}) > 0$ – a contradiction. Suppose, instead, that the deviator does make positive sales of the 90 policy (which requires that $P'_{90} \leq \underline{AC}_{90}$). Since it is a profitable deviation it must also make positive sales of the 60 policy (since $P'_{90} \leq \underline{AC}_{90}$, it can't make a positive profit selling only the 90 policy). But, in this case $P'_{90} \leq \underline{AC}_{90} < AC_{90}(\Delta P')$, so we must have $\Pi_{60}(P'_{90}, P'_{60}) > 0$ since it is a profitable deviation. In turn, letting $\widehat{P}_{60} \equiv \underline{AC}_{90} - \Delta P' \geq P'_{60}$, this implies that $\Pi_{60}(\underline{AC}_{90}, \widehat{P}_{60}) \geq \Pi_{60}(P'_{90}, P'_{60}) > 0$ – a contradiction. □

Lemma 12. *Among all price pairs (P_{90}, P_{60}) at which both policies break even and there are positive sales of the 60 policy, only the one with the lowest sales of the 60 policy (i.e., having $\Delta P = \underline{\Delta P}^{BE}$) can be a Riley equilibrium.*

Proof. Suppose there are two price pairs that are Riley Equilibria $P' = (P'_{90}, P'_{60})$ and $P'' = (P''_{90}, P''_{60})$ with both having positive and differing 60 shares. We have already seen in the proof of Lemma 2 that, starting at P'', a single-policy deviation to P'_{90} is profitable. Now, observe that the lowest safe 60 price given P'_{90} is P'_{60}; i.e., $\underline{P}_{60}(P'_{90}) = P'_{60}$, so $\Pi_{90}(P'_{90}, \underline{P}_{60}(P'_{90})) = 0$. Hence, there are no safe reactions that make the deviator incur a loss (Remark 3). This implies that (P''_{90}, P''_{60}) is not a Riley equilibrium, which is a contradiction. □

Lemma 13. *Suppose that at $P^* = (P^*_{90}, P^*_{60})$ there are positive sales of the 60 policy (so $\Delta P^* \in (\underline{\theta}, \bar{\theta}]$) and both policies break even. Then P^* is a Riley equilibrium if and only if there are no single-policy Riley profitable deviations in P_{90}.*

Proof. Consider a multi-policy profitable Riley deviation $P' = (P'_{90}, P'_{60})$. We will show that we necessarily have $\Pi_{90}(P'_{90}, P^*_{60}) > 0$ and $\Pi_{90}(P'_{90}, \widetilde{P}_{60}) \geq 0$ for all $\widetilde{P}_{60} \in [\underline{P}_{60}(P'_{90}), P^*_{60}]$. Thus, a single-policy deviation to P'_{90} would be a profitable Riley deviation. This will imply (see Remark 3) that in looking for Riley profitable deviations when the 60 policy has positive sales, we can restrict attention to single-policy deviations in P_{90}.

The claim is immediate if $P'_{60} > P^*_{60}$ since then dropping offer P'_{60} would affect neither the deviation profit, nor the deviator's profit after any reaction. So henceforth we shall assume that $P'_{60} \leq P^*_{60}$. Moreover, we must have $P'_{90} \leq P^*_{90}$: otherwise the deviator can sell only the 60 policy at price $P'_{60} \leq P^*_{60} = AC_{60}(\Delta P^*) \leq AC_{60}(\Delta P')$, contradicting P' being a profitable Riley deviation. So $P' \leq P^*$.

Next, observe that we must have $\Delta P' < \Delta P^*$ and an increased share of the 90 policy being purchased. If not, then since the average costs of both policies would be no lower than they were before the deviation, and both deviation prices would be weakly lower, the deviation could not generate a strictly positive profit. Note that this also implies that we must have $P'_{90} < P^*_{90}$.

Suppose, first, that $\underline{P}_{60}(P'_{90}) < P'_{60}$. If $\Pi_{90}(P'_{90}, \underline{P}_{60}(P'_{90})) < 0$, then the safe single-policy reaction to $\underline{P}_{60}(P'_{90})$ makes the deviator incur losses, in contradiction to the assumption that P' is a profitable Riley deviation. So in this case we must have $\Pi_{90}(P'_{90}, \underline{P}_{60}(P'_{90})) \geq 0$. Moreover, there must be positive sales of the 90 policy at prices $(P'_{90}, \underline{P}_{60}(P'_{90}))$ because, if not, then (see Remark 1) $\underline{P}_{60}(P'_{90}) = \overline{AC}_{60} \geq P^*_{60} \geq P'_{60}$. Thus, $\Pi_{90}(P'_{90}, \widetilde{P}_{60}) > 0$ for all $\widetilde{P}_{60} \in (\underline{P}_{60}(P'_{90}), P^*_{60}]$, implying that the single-policy deviation to P'_{90} is a profitable Riley deviation.

On the other hand, if $\underline{P}_{60}(P'_{90}) \geq P'_{60}$, then $\Pi_{60}(P'_{90}, P'_{60}) \leq 0$, which implies that $\Pi_{90}(P'_{90}, P'_{60}) > 0$ (since the deviation to P' is profitable). This, in turn, implies that $\Pi_{90}(P'_{90}, P^*_{60}) > 0$, which establishes the result. □

Lemma 14. *A Riley equilibrium exists.*

Proof. Suppose otherwise. It is immediate that if at prices $(P_{90}, P_{60}) = (\underline{AC}_{90}, \overline{AC}_{60})$ all consumers buy policy 60, then by Lemma 13 this is a Riley equilibrium as there is no single-policy deviation in P_{90} that can earn a strictly positive profit, yielding a contradiction. Thus, we henceforth assume that there are positive sales of the 90 policy at $(P_{90}, P_{60}) = (\underline{AC}_{90}, \overline{AC}_{60})$.

Given Lemma 12, we need to show that either a price configuration leading to all-in-90 or the break-even price configuration with $\Delta P = \underline{\Delta P}^{BE}$ is a Riley Equilibrium.

Let $P^{**} = (P^{**}_{90}, P^{**}_{60})$ be the break-even price configuration with $\Delta P^{**} = \underline{\Delta P}^{BE} > \underline{\theta}$. Note also that $P^{**}_{60} = \underline{P}_{60}(P^{**}_{90})$. [This is immediate if $\Delta P^{**} \equiv P^{**}_{90} - P^{**}_{60} \in (\underline{\theta}, \overline{\theta})$. If $\Delta P^{**} = \overline{\theta}$, then (by Lemma 10) $P^{**}_{60} = \overline{AC}_{60}$ and $P^{**}_{90} = \overline{AC}_{60} + \overline{\theta}$. By Remark 1, the latter equality implies that $\underline{P}_{60}(P^{**}_{90}) = \overline{AC}_{60}$.]

If P^{**} is not a Riley Equilibrium, by Lemma 13 there is a single-policy deviation $P'_{90} < P^{**}_{90}$ with $\Pi_{90}(P'_{90}, P^{**}_{60}) > 0$ that does not incur losses when facing the reaction of the lowest safe price $\underline{P}_{60}(P'_{90}) \leq \underline{P}_{60}(P^{**}_{90}) = P^{**}_{60}$, so that $\Pi_{90}(P'_{90}, \underline{P}_{60}(P'_{90})) \geq 0$. Note also that there must be positive sales of the 90 policy at prices $P' = (P'_{90}, \underline{P}_{60}(P'_{90}))$ since otherwise (by Remark 1) we have $\underline{P}_{60}(P'_{90}) = \overline{AC}_{60} \geq P^{**}_{60}$, which would imply that there are no sales of the 90 policy at prices (P'_{90}, P^{**}_{60}), contradicting $\Pi_{90}(P'_{90}, P^{**}_{60}) > 0$.

Now, suppose that everyone-in-90 is not a Riley Equilibrium. Then by Lemma 11 there is a $\widehat{P}_{60}$ such that $\Pi_{60}(\underline{AC}_{90}, \widehat{P}_{60}) > 0$, implying that $\underline{P}_{60}(\underline{AC}_{90}) < \widehat{P}_{60}$ and at prices $(\underline{AC}_{90}, \underline{P}_{60}(\underline{AC}_{90}))$ there are positive sales of the 60 policy. Moreover, there must also be positive sales of the 90 policy at prices $(\underline{AC}_{90}, \underline{P}_{60}(\underline{AC}_{90}))$: if not, by Remark 1 we would have $\underline{P}_{60}(\underline{AC}_{90}) = \overline{AC}_{60}$, but this contradicts (recall the first paragraph of the proof) that there are positive sales of the 90 policy at $(P_{90}, P_{60}) = (\underline{AC}_{90}, \overline{AC}_{60})$. Because there are positive sales of both policies at prices $(P_{90}, P_{60}) = (\underline{AC}_{90}, \underline{P}_{60}(\underline{AC}_{90}))$, we have $\Pi_{90}(\underline{AC}_{90}, \underline{P}_{60}(\underline{AC}_{90})) < 0$. Continuity of the function $\Pi_{90}(P_{90}, \underline{P}_{60}(P_{90}))$ in P_{90} then implies that there is a $\widetilde{P}_{90} \in (\underline{AC}_{90}, P'_{90}]$ at which $\Pi_{90}(\widetilde{P}_{90}, \underline{P}_{60}(\widetilde{P}_{90})) = 0$. Hence, both policies break even at price pair $\widetilde{P} = (\widetilde{P}_{90}, \underline{P}_{60}(\widetilde{P}_{90}))$.

Finally, we establish the result by showing that there are positive sales of the 60 policy at prices $\widetilde{P}$, and that $\Delta\widetilde{P} < \Delta P^{**}$, which contradicts P^{**} being the price pair with the lowest ΔP among those with positive sales of the 60 policy. For the first point, note that because there are positive sales of the 60 policy at prices $(\underline{AC}_{90}, \underline{P}_{60}(\underline{AC}_{90}))$, Remark 1 tells us that there are positive sales of the 60 policy at price pair $\widetilde{P}$ (and also at prices P'). For the second point, observe that because there are positive sales of both policies at prices P', by Remark 1 we have $\Delta\widetilde{P} \leq \Delta P' < \Delta P^{**}$. This establishes the result. $\square$

A.3 Wilson Equilibria

Recall that a price configuration $P = (P_{90}, P_{60})$ is a Wilson equilibrium (WE) if there is no deviation by an entrant to a price pair that is strictly profitable once any offers are withdrawn that make losses after the deviation.[55] We will say that a deviation from price configuration P that is strictly profitable after any such withdrawals is a "profitable Wilson deviation." Note that no policy 60 offers will ever be withdrawn after a deviation, because a reduction in P_{90} can never cause a P_{60} offer to make losses (since a reduction in P_{90} lowers AC_{60}).

We characterize Wilson equilibria (WE) through a series of lemmas that imply Proposition 3. First, we identify some properties that any WE must satisfy:

Lemma 15. *If $P^w = (P^w_{90}, P^w_{60})$ is a WE price configuration, then*

(a) $\Pi(P^w_{90}, P^w_{60}) = 0$;

(b) $\Pi_{90}(P'_{90}, P^w_{60}) \leq 0$ *for all* $P'_{90} \leq P^w_{90}$*;*

(c) $\Delta P^w = (P^w_{90} - P^w_{60}) \leq \underline{\Delta P}^{BE}$*, the lowest break-even* ΔP *with positive sales of the 60 policy.*

Proof. (a) If $\Pi(P^w_{90}, P^w_{60}) < 0$, then some firm would be better off dropping its offers, while if $\Pi(P^w_{90}, P^w_{60}) > 0$ then an entrant could profit by offering $(P^w_{90} - \varepsilon, P^w_{60} - \varepsilon)$ for sufficiently small $\varepsilon > 0$. (b) If this is violated at P'_{90}, then $\Pi_{90}(P'_{90} - \varepsilon, P^w_{60}) > 0$ for sufficiently small $\varepsilon > 0$. A entrants' offering of $P'_{90} - \varepsilon$ would

[55] Note that since at least one of P_{90} and P_{60} is undercut by any profitable entrant deviation, there is no ambiguity about which polices to withdraw in the event that one of the offers in the price configuration makes losses.

be a profitable Wilson deviation. (c) This is immediate if $\underline{\Delta P}^{BE} = \overline{\theta}$. So suppose that $\underline{\Delta P}^{BE} < \overline{\theta}$ and that $\Delta P^w > \underline{\Delta P}^{BE}$. Since both policies break even at $\underline{\Delta P}^{BE}$, and $\Pi_{60}(P_{90}^w, P_{60}^w) \geq 0$ by parts (a) and (b), it must be that the break-even price configuration associated with $\underline{\Delta P}^{BE}$, $(\underline{P}_{90}^{BE}, \underline{P}_{60}^{BE})$, has $\underline{P}_{60}^{BE} = AC_{60}(\underline{\Delta P}^{BE}) < AC_{60}(\Delta P^w) \leq P_{60}^w$. Since $\underline{P}_{60}^{BE} < P_{60}^w$ and $\underline{\Delta P}^{BE} < \Delta P^w$, we also have $\underline{P}_{90}^{BE} < P_{90}^w$. So an entrant's offer of $(\underline{P}_{90}^{BE} + \varepsilon, \underline{P}_{60}^{BE} + \varepsilon)$ for sufficiently small $\varepsilon > 0$ is a profitable Wilson deviation. $\square$

Consider the following problem:

$$\begin{array}{cc} \min_{(P_{90},P_{60})} & \min P_{60} \\ \text{s.t.} & \text{(i) } \Pi(P_{90}, P_{60}) = 0 \\ & \text{(ii) } \Pi_{90}(P'_{90}, P_{60}) \leq 0 \text{ for all } P'_{90} \leq P_{90} \\ & \text{(iii) } P_{90} - P_{60} \in [\underline{\theta}, \underline{\Delta P}^{BE}] \end{array} \tag{10}$$

Lemma 16. *Any $P^* = (P_{90}^*, P_{60}^*)$ that solves problem (10) is a WE price configuration.*

Proof. We construct an equilibrium in which all prices $P \geq P^*$ are offered by multiple firms and each firm has an equal share of sales of both policies. Thus, all active firms earn zero, and we need only consider deviations by entrants.

To begin, it follows from constraint (ii) of problem (10), and the fact that 60 offers are never withdrawn, that there is no profitable Wilson deviation in which an entrant makes sales only of the 90 policy (which would require a price $\widehat{P}_{90} < P_{90}^*$).

Next, there is no profitable Wilson deviation in which an entrant makes sales only of the 60 policy. Suppose there were and let the deviation 60 price be $\widehat{P}_{60} < P_{60}^*$. If everyone buys the 60 policy at prices $(P_{90}^*, \widehat{P}_{60})$ then no 90 policy offers will be withdrawn and $\widehat{P}_{60} > \overline{AC}_{60}$. But then prices $(P_{90}^*, \overline{AC}_{60})$ would be feasible in problem (10) and attain a lower value of P_{60} than P_{60}^*, contradicting P^* being a solution. Suppose instead that some consumers still buy the 90 policy at prices $(P_{90}^*, \widehat{P}_{60})$. Then $\Pi_{90}(P_{90}^*, \widehat{P}_{60}) < 0$, which implies that offer P_{90}^* will be withdrawn, as will every P_{90} up to the lowest $\overline{P}_{90}$ above P_{90}^* such that $\Pi_{90}(\overline{P}_{90}, \widehat{P}_{60}) = 0$. The entrant's profit will therefore be $\Pi_{60}(\overline{P}_{90}, \widehat{P}_{60})$. However, it cannot be that $\Pi_{60}(\overline{P}_{90}, \widehat{P}_{60}) > 0$: if so then we have $\Pi(\overline{P}_{90}, \widehat{P}_{60}) > 0$. But this would imply that there is an $\varepsilon > 0$ such that price pair $(\overline{P}_{90} - \varepsilon, \widehat{P}_{60} - \varepsilon)$ is feasible in problem (10) and achieves a lower P_{60} than P_{60}^*, a contradiction to P^* solving problem (10).[56]

Finally, suppose that there is a profitable Wilson deviation for an entrant offering $\widehat{P} = (\widehat{P}_{90}, \widehat{P}_{60})$, in which the entrant makes sales of both policies. Then since offers for the 60 policy are never withdrawn, $\widehat{P}_{60} < P_{60}^*$. We first argue that $\Pi_{90}(P_{90}, \widehat{P}_{60}) \leq 0$ for all $P_{90} \leq \widehat{P}_{90}$. If $\widehat{P}_{90} < P_{90}^*$, then this follows because P^* satisfies constraint (ii) and $\widehat{P}_{60} < P_{60}^*$. If, instead, $\widehat{P}_{90} > P_{90}^*$, then it follows because the entrant can make sales of the 90 policy only if $\Pi_{90}(P_{90}, \widehat{P}_{60}) < 0$ for all $P_{90} < \widehat{P}_{90}$, so that rivals' offers are withdrawn. Next, observe that if $\Pi_{90}(\widehat{P}_{90}, \widehat{P}_{60}) \leq 0$ and $\Pi(\widehat{P}_{90}, \widehat{P}_{60}) > 0$, then for some $\varepsilon > 0$ price

[56] This ε would set $\Pi(\overline{P}_{90} - \varepsilon, \widehat{P}_{60} - \varepsilon) = 0$, and would satisfy constraint (ii) of problem (10) since $\Pi_{90}(P_{90}, \widehat{P}_{60} - \varepsilon) \leq 0$ for all $P_{90} \leq \overline{P}_{90}$.

pair $(\widehat{P}_{90}-\varepsilon, \widehat{P}_{60}-\varepsilon)$ is feasible in problem (10) and achieves a lower P_{60} than P^*_{60}, a contradiction to P^* solving problem (10).[57] $\square$

To solve for the Wilson equilibrium, we examine a relaxed version of problem (10). For $\Delta P \in [\underline{\theta}, \overline{\theta}]$, we first define $P^{BE}_{60}(\Delta P)$ by

$$[P^{BE}_{60}(\Delta P) - AC_{60}(\Delta P)]F(\Delta P) + [P^{BE}_{60}(\Delta P) + \Delta P - AC_{90}(\Delta P)][1-F(\Delta P)] = 0,$$

and $P^{BE}_{90}(\Delta P) \equiv P^{BE}_{60}(\Delta P) + \Delta P$. Note that $P^{BE}_{60}(\Delta P)$ and $P^{BE}_{90}(\Delta P)$ are continuous functions. Note as well that, for $\Delta P \in [\underline{\theta}, \overline{\theta}]$, $[P^{BE}_{60}(\Delta P) - AC_{60}(\Delta P)] \gtreqless 0$ if and only if $\Delta AC(\Delta P) \gtreqless \Delta P$.[58]

We will consider the relaxed problem

$$\min_{\Delta P \in [\underline{\theta}, \underline{\Delta P}^{BE}]} \quad P^{BE}_{60}(\Delta P) \tag{11}$$

Note that in problem (11) the constraint set is closed and bounded, and the objective function is continuous, so a solution exists. We also note the equivalence of this problem to the problem of finding the profit-maximizing multi-policy Nash deviation from price configuration $(P^{BE}_{90}(\underline{\Delta P}^{BE}), P^{BE}_{60}(\underline{\Delta P}^{BE}))$:

$$\max_{\Delta P \in [\underline{\theta}, \underline{\Delta P}^{BE}]} \quad \Pi(P^{BE}_{60}(\underline{\Delta P}^{BE}) + \Delta P, P^{BE}_{60}(\underline{\Delta P}^{BE})) \tag{12}$$

Lemma 17. *$Arg\min_{\Delta P \in [\underline{\theta}, \underline{\Delta P}^{BE}]} P^{BE}_{60}(\Delta P) = Arg\max_{\Delta P \in [\underline{\theta}, \underline{\Delta P}^{BE}]} \Pi(P^{BE}_{60}(\underline{\Delta P}^{BE}) + \Delta P, P^{BE}_{60}(\underline{\Delta P}^{BE}))$.*

Proof. Letting $\delta(\Delta P) \equiv P^{BE}_{60}(\underline{\Delta P}^{BE}) - P^{BE}_{60}(\Delta P)$, we have

$$\begin{aligned}
\Pi(P^{BE}_{60}(\underline{\Delta P}^{BE}) + \Delta P, P^{BE}_{60}(\underline{\Delta P}^{BE})) &= \Pi(P^{BE}_{60}(\Delta P) + \Delta P + \delta, P^{BE}_{60}(\Delta P) + \delta) \\
&= \Pi(P^{BE}_{60}(\Delta P) + \Delta P, P^{BE}_{60}(\Delta P)) + \delta \\
&= P^{BE}_{60}(\underline{\Delta P}^{BE}) - P^{BE}_{60}(\Delta P),
\end{aligned}$$

so for any ΔP and $\Delta P'$ we have

$$\Pi(P^{BE}_{60}(\underline{\Delta P}^{BE}) + \Delta P, P^{BE}_{60}(\underline{\Delta P}^{BE})) - \Pi(P^{BE}_{60}(\underline{\Delta P}^{BE}) + \Delta P', P^{BE}_{60}(\underline{\Delta P}^{BE})) = P^{BE}_{60}(\Delta P') - P^{BE}_{60}(\Delta P).$$

$\square$

Thus, the solution to the relaxed problem (11) is exactly the $\Delta P \leq \underline{\Delta P}^{BE}$ that maximizes the multi-policy deviation profits from $\underline{\Delta P}^{BE}$. The usefulness of the relaxed problem stems from the following result [whose assumption that $\Delta AC(\underline{\theta}) > \underline{\theta}$ is satisfied in our data]:

[57] This ε would set $\Pi(\widehat{P}_{90}-\varepsilon, \widehat{P}_{60}-\varepsilon) = 0$, and would satisfy constraint (ii) of problem (10) since $\Pi_{90}(P_{90}, \widehat{P}_{60}-\varepsilon) \leq 0$ for all $P_{90} \leq \widehat{P}_{90}$.

[58] This follows since

$$\Delta AC(\Delta P) \gtreqless \Delta P \Leftrightarrow P^{BE}_{60}(\Delta P) - AC_{60}(\Delta P) \gtreqless P^{BE}_{90}(\Delta P) - AC_{90}(\Delta P).$$

Lemma 18. *Suppose that $\Delta AC(\underline{\theta}) > \underline{\theta}$ and that $\Delta P^* = \arg\min_{\Delta P \in [\underline{\theta}, \underline{\Delta P}^{BE}]} P_{60}^{BE}(\Delta P)$. Then the price configuration $(P_{90}^{BE}(\Delta P^*), P_{60}^{BE}(\Delta P^*))$ is the unique solution to problem (10).*

Proof. By construction $(P_{90}^{BE}(\Delta P), P_{60}^{BE}(\Delta P^*))$ satisfies constraints (i) and (iii) of problem (10). We therefore need only show that $(P_{90}^{BE}(\Delta P^*), P_{60}^{BE}(\Delta P^*))$ satisfies constraint (ii). Observe that when $\Delta AC(\underline{\theta}) > \underline{\theta}$, at any $\Delta P \in [\underline{\theta}, \underline{\Delta P}^{BE}]$ we have $\Delta AC(\Delta P) > \Delta P$. This implies that for all $\Delta P \in [\underline{\theta}, \underline{\Delta P}^{BE}]$

$$\begin{aligned} 0 &\geq \Pi_{90}(P_{90}^{BE}(\Delta P), P_{60}^{BE}(\Delta P)) \\ &\geq \Pi_{90}(P_{90}^{BE}(\Delta P), P_{60}^{BE}(\Delta P^*)), \end{aligned}$$

where the second inequality follows because $P_{60}^{BE}(\Delta P^*) \leq P_{60}^{BE}(\Delta P)$ for all $\Delta P \in [\underline{\theta}, \underline{\Delta P}^{BE}]$ by virtue of ΔP^* being the solution to problem (11). Continuity of $P_{90}^{BE}(\Delta P)$ in ΔP then implies that

$$0 \geq \Pi_{90}(P_{90}, P_{60}^{BE}(\Delta P^*)) \text{ for all } P_{90} \in [\underline{AC}_{90}, P_{90}^{BE}(\Delta P^*)].$$

Since we also have that

$$0 \geq \Pi_{90}(P_{90}, P_{60}^{BE}(\Delta P^*)) \text{ for all } P_{90} \leq \underline{AC}_{90},$$

$(P_{90}^{BE}(\Delta P^*), P_{60}^{BE}(\Delta P^*))$ satisfies constraint (ii) of problem (10). □

Finally, we show that the solution P^* to problem (10) is the *only* WE whenever $\Delta P^* \in (\underline{\theta}, \underline{\Delta P}^{BE})$.

Lemma 19. *Suppose that there is a unique solution P^* of problem (10) and that $\Delta P^* \in (\underline{\theta}, \underline{\Delta P}^{BE})$. Then P^* is the unique WE price configuration.*[59]

Proof. Lemma 15 shows that any WE price configuration must satisfy the constraints of problem (10). We next argue that any price configuration $\widetilde{P} = (\widetilde{P}_{90}, \widetilde{P}_{60})$ that satisfies the constraints but is not a solution cannot be a WE price configuration. By definition, $P_{60}^* < \widetilde{P}_{60}$.

If $(P_{90}^*, P_{60}^*) << (\widetilde{P}_{90}, \widetilde{P}_{60})$ then at price configuration $(\widetilde{P}_{90}, \widetilde{P}_{60})$ an entrant has a profitable Wilson deviation to $(P_{90}^* + \varepsilon, P_{60}^* + \varepsilon)$ for small $\varepsilon > 0$. So, for the rest of the proof, suppose instead that $P_{90}^* \geq \widetilde{P}_{90}$, which also implies that $\Delta\widetilde{P} < \Delta P^*$ since $P_{60}^* < \widetilde{P}_{60}$.

Observe, first, that we must then have $\underline{AC}_{90} - \underline{\theta} > \underline{AC}_{60}$, since $P_{60}^* > \underline{AC}_{60}$ [which follows from there being sales of the 60 policy at P^* and $\Pi_{60}(P_{90}^*, P_{60}^*) \geq 0$] and $(\underline{AC}_{90}, \underline{AC}_{90} - \underline{\theta})$ is feasible in problem (10). Thus, $\Delta AC(\underline{\theta}) > \underline{\theta}$, which implies that $\Delta AC(\Delta P) > \Delta P$ for all $\Delta P \in (\underline{\theta}, \underline{\Delta P}^{BE})$ and, in turn, that $\Pi_{90}(P_{90}^{BE}(\Delta P), P_{60}^{BE}(\Delta P)) < 0$ for all $\Delta P \in (\underline{\theta}, \underline{\Delta P}^{BE})$. Moreover, continuity implies that for each $P_{90} \in [\widetilde{P}_{90}, P_{90}^*]$, there is a $\Delta P \in (\underline{\theta}, \underline{\Delta P}^{BE})$ such that $P_{90}^{BE}(\Delta P) = P_{90}$. Thus, since $P_{60}^* < P_{60}^{BE}(\Delta P)$ for all ΔP and there are positive sales of the 90 policy at P^*, we have $\Pi_{90}(P_{90}, P_{60}^*) < 0$ for all $P_{90} \in (\widetilde{P}_{90}, P_{90}^*]$ and at $\widetilde{P}_{90}$ if $\widetilde{P}_{90} > \underline{AC}_{90}$.

We will consider two cases:

[59] We conjecture, but have not proven, that the result extends to cases in which $\Delta P^* = \underline{\Delta P}^{BE}$.

(i) $P_{90}^* = \widetilde{P}_{90}$. In this case, there are sales of the 90 policy at prices $\widetilde{P}$, and we have $0 \geq \Pi_{90}(\widetilde{P}_{90}, \widetilde{P}_{60}) = \Pi_{90}(P_{90}^*, \widetilde{P}_{60}) > \Pi_{90}(P_{90}^*, P_{60}^*) = \Pi_{90}(\widetilde{P}_{90}, P_{60}^*)$. So at $\widetilde{P}$ an entrant has a profitable Wilson deviation offering prices $(P_{90}^* + \varepsilon, P_{60}^* + \varepsilon)$ for $\varepsilon > 0$ such that $\Pi_{90}(P_{90}', P_{60}^* + \varepsilon) < 0$ for all $P_{90}' \in [\widetilde{P}_{90}, P_{90}^* + \varepsilon]$ and $P_{60}^* + \varepsilon < \widetilde{P}_{60}$, which results in all 90 policy offers in $[\widetilde{P}_{90}, P_{90}^* + \varepsilon]$ being withdrawn.

(ii) $P_{90}^* > \widetilde{P}_{90}$. Note first that if $\Delta\widetilde{P} = \underline{\theta}$ (so there are no sales of the 60 policy at $\widetilde{P}$) then $\widetilde{P}_{90} = \underline{AC}_{90}$. There must then be positive sales of the 60 policy at $(\underline{AC}_{90}, P_{60}^*)$ for otherwise $(\underline{AC}_{90}, P_{60}^*)$ would satisfy constraints (i) and (ii) of problem (10), contradicting P^* being the unique solution. This implies that $\Pi_{90}(\underline{AC}_{90}, P_{60}^*) < 0$. As well, recall that we have $\Pi_{90}(P_{90}, P_{60}^*) < 0$ for all $P_{90} \in (\widetilde{P}_{90}, P_{90}^*]$ and at $\widetilde{P}_{90}$ if $\widetilde{P}_{90} > \underline{AC}_{90}$. Thus, in this case, an entrant again has a profitable Wilson deviation offering prices $(P_{90}^* + \varepsilon, P_{60}^* + \varepsilon)$ for $\varepsilon > 0$ such that $\Pi_{90}(P_{90}', P_{60}^* + \varepsilon) < 0$ for all $P_{90}' \in [\widetilde{P}_{90}, P_{90}^* + \varepsilon]$ and $P_{60}^* + \varepsilon < \widetilde{P}_{60}$, which results in all 90 policy offers in in $[\widetilde{P}_{90}, P_{90}^* + \varepsilon]$ being withdrawn. $\square$

B Appendix: Cost Model Setup and Estimation

This appendix describes the details of the cost model, which is summarized at a high-level in section 3, and similar to that used in Handel (2013). The output of this model, F_{jkt}, is a family-plan-time-specific distribution of predicted out-of-pocket expenditures for the upcoming year. This distribution is an important input into the empirical choice model, where it enters as a family's predictions of its out-of-pocket expenses at the time of plan choice, for each plan option. We predict this distribution in a sophisticated manner that incorporates (i) past diagnostic information (ICD-9 codes) (ii) the Johns Hopkins ACG predictive medical software package (iii) a non-parametric model linking modeled health risk to total medical expenditures using observed cost data and (iv) a detailed division of medical claims and health plan characteristics to precisely map total medical expenditures to out-of-pocket expenses. The level of precision we gain from the cost model leads to more credible estimates of the choice parameters of primary interest (e.g., risk preferences and health risk). Crucially, the cost model output is also used to predict consumer expected average costs for the upcoming year, λ, which is used to determine plan costs (as a function of who selects which plans) in our equilibrium analyses.

In order to predict expenses in a precise manner, we categorize the universe of total medical claims into four mutually exclusive and exhaustive subdivisions of claims using the claims data. These categories are (i) hospital and physician services (ii) pharmacy (iii) mental health and (iv) physician office visits. We divide claims into these four specific categories so that we can accurately characterize the plan-specific mappings from total claims to out-of-pocket expenditures since each of these categories maps to out-of-pocket expenditures in a different manner. We denote this four dimensional vector of claims $\mathbf{C}_{it}$ and any given element of that vector $C_{d,it}$ where $d \in D$ represents one of the four categories and i denotes an individual (employee or dependent). After describing how we predict this vector of claims for a given individual, we return to the question of how we determine out-of-pocket expenditures in plan k given $\mathbf{C}_{it}$.

Denote an individual's past year of medical diagnoses and payments by ξ_{it} and the demographics age and sex by ζ_{it}. We use the ACG software mapping, denoted A, to map these characteristics into a predicted mean level of health expenditures for the upcoming year, denoted θ:

$$A : \xi \times \zeta \rightarrow \theta$$

In addition to forecasting a mean level of total expenditures, the software has an application that predicts future mean pharmacy expenditures. This mapping is analogous to A and outputs a prediction κ for future pharmacy expenses.

We use the predictions θ and κ to categorize similar groups of individuals across each of four claims categories in vector in $\mathbf{C}_{it}$. Then for each group of individuals in each claims category, we use the actual ex post realized claims for that group to estimate the ex ante distribution for each individual under the assumption that this distribution is identical for all individuals within the cell. Individuals are categorized into cells based on different metrics for each of the four elements of $\mathbf{C}$:

$$\begin{aligned} \text{Pharmacy:} &\quad \kappa_{it} \\ \text{Hospital / Physician (Non-OV):} &\quad \theta_{it} \\ \text{Physician Office Visit:} &\quad \theta_{it} \\ \text{Mental Health:} &\quad C_{MH,i,t-1} \end{aligned}$$

For pharmacy claims, individuals are grouped into cells based on the predicted future mean pharmacy claims measure output by the ACG software, κ_{it}. For the categories of hospital / physician services (non office visit) and physician office visit claims individuals are grouped based on their mean predicted total future health expenses, θ_{it}. Finally, for mental health claims, individuals are grouped into categories based on their mental health claims from the previous year, $C_{MH,i,t-1}$ since (i) mental health claims are very persistent over time in the data and (ii) mental health claims are generally uncorrelated with other health expenditures in the data. For each category we group individuals into a number of cells between 8 and 10, taking into account the tradeoff between cell size and precision. The minimum number of individuals in any cell is 73 while almost all cells have over 500 members. Thus, since there are four categories of claims, each individual can belong to one of approximately 10^4 or 10,000 combination of cells.

Denote an arbitrary cell within a given category d by z. Denote the population in a given category-cell combination (d, z) by I_{dz}. Denote the empirical distribution of ex-post claims in this category for this population $\hat{G_{I_{dz}}}(\cdot)$. Then we assume that each individual in this cell has a distribution equal to a continuous fit of $\hat{G_{I_{dz}}}(\cdot)$, which we denote G_{dz}:

$$\varpi : \hat{G_{I_{dz}}}(\cdot) \to G_{dz}$$

We model this distribution continuously in order to easily incorporate correlations across d. Otherwise, it would be appropriate to use $G_{I_{dz}}$ as the distribution for each cell.

The above process generates a distribution of claims for each d and z but does not model correlation over D. It is important to model correlation across claims categories because it is likely that someone with a bad expenditure shock in one category (e.g., hospital) will have high expenses in another area (e.g., pharmacy). We model correlation at the individual level by combining marginal distributions G_{idt} $\forall$ d with empirical data on the rank correlations between pairs (d, d').[60] Here, G_{idt} is the distribution G_{dz} where $i \in I_{dz}$ at time t. Since correlations are modeled across d we pick the metric θ to group people into cells for the basis of determining correlations (we use the same cells that we use to determine group people for hospital and physician office visit claims). Denote these cells based on θ by z_θ. Then for each cell z_θ denote the empirical rank correlation between claims of type d and type d' by $\rho_{z_\theta}(d, d')$.

[60] It is important to use rank correlations here to properly combine these marginal distribution into a joint distribution. Linear correlation would not translate empirical correlations to this joint distribution appropriately.

Then, for a given individual i we determine the joint distribution of claims across D for year t, denoted $H_{it}(\cdot)$, by combining i's marginal distributions for all d at t using $\rho_{z_\theta}(d, d')$:

$$\Psi : G_{iDt} \times \rho_{z_{\theta_{it}}}(D, D') \rightarrow H_{it}$$

Here, G_{iDt} refers to the set of marginal distributions $G_{idt} \forall d \in D$ and $\rho_{z_{\theta_{it}}}(D, D')$ is the set of all pairwise correlations $\rho_{z_{\theta_{it}}}(d, d') \forall (d, d'^2$. In estimation we perform Ψ by using a Gaussian copula to combine the marginal distribution with the rank correlations, a process which we describe momentarily.

The final part of the cost model maps the joint distribution H_{it} of the vector of total claims C over the four categories into a distribution of out of pocket expenditures for each plan. For each of the three plan options we construct a mapping from the vector of claims C to out-of-pocket expenditures X_k:

$$\Omega_k : C \rightarrow X_k$$

This mapping takes a given draw of claims from H_{it} and converts it into the out-of-pocket expenditures an individual would have for those claims in plan k. This mapping accounts for plan-specific features such as the deductible, co-insurance, co-payments, and out-of-pocket maximums described in the text. We test the mapping Ω_k on the actual realizations of the claims vector C to verify that our mapping comes close to reconstructing the true mapping. Our mapping is necessarily simpler and omits things like emergency room co-payments and out of network claims. We constructed our mapping with and without these omitted categories to insure they did not lead to an incremental increase in precision. We find that our categorization of claims into the four categories in C passed through our mapping Ω_k closely approximates the true mapping from claims to out-of-pocket expenses. Further, we find that it is important to model all four categories described above: removing any of the four makes Ω_k less accurate. See Handel (2013) for figures describing this validation exercise with the data used in this paper.

Once we have a draw of X_{ikt} for each i (claim draw from H_{it} passed through Ω_k) we map individual out-of-pocket expenditures into family out-of-pocket expenditures. For families with less than two members this involves adding up all the within family X_{ikt}. For families with more than three members there are family level restrictions on deductible paid and out-of-pocket maximums that we adjust for. Define a family j as a collection of individuals i_j and the set of families as J. Then for a given family out-of-pocket expenditures are generated:

$$\Gamma_k : X_{i_j,kt} \rightarrow X_{jkt}$$

To create the final object of interest, the family-plan-time specific distribution of out of pocket expenditures $F_{jkt}(\cdot)$, we pass the claims distributions H_{it} through Ω_k and combine families through Γ_k. $F_{jkt}(\cdot)$ is then used as an input into the choice model that represents each family's information set over future medical expenses at the time of plan choice. Eventually, we also use H_{it} to calculate total plan cost when we analyze counterfactual plan pricing based on the average cost of enrollees.

We note that the decision to do the cost model by grouping individuals into cells, rather then by specifying a more continuous form, has costs and benefits. The cost is that all individuals within a given cell for a given type of claims are treated identically. The benefit is that our method produces local cost estimates for each individual that are not impacted by the combination of functional form and the health risk of medically different individuals. Also, the method we use allows for flexible modeling across claims categories. Finally, we note that we map the empirical distribution of claims to a continuous representation because this is convenient for building in correlations in the next step. The continuous distributions we generate very closely fit the actual empirical distribution of claims across these four categories.

Cost Model Identification and Estimation. The cost model is identified based on the two assumptions of (i) no moral hazard / selection based on private information and (ii) that individuals within the same cells for claims d have the same ex ante distribution of total claims in that category. Once these assumptions are made, the model uses the detailed medical data, the Johns Hopkins predictive algorithm, and the plan-specific mappings for out of pocket expenditures to generate the final output $F_{jkt}(\cdot)$. These assumptions, and corresponding robustness analyses, are discussed at more length in the main text and in Handel (2013).

Once we group individuals into cells for each of the four claims categories, there are two statistical components to estimation. First, we need to generate the continuous marginal distribution of claims for each cell z in claim category d, G_{dz}. To do this, we fit the empirical distribution of claims $G_{I_{dz}}$ to a Weibull distribution with a mass of values at 0. We use the Weibull distribution instead of the lognormal distribution, which is traditionally used to model medical expenditures, because we find that the lognormal distribution overpredicts large claims in the data while the Weibull does not. For each d and z the claims greater than zero are estimated with a maximum likelihood fit to the Weibull distribution:

$$\max_{(\alpha_{dz},\beta_{dz})} \Pi_{i\in I_{dz}} \frac{\beta_{dz}}{\alpha_{dz}}\left(\frac{c_{id}}{\alpha_{dz}}\right)^{\beta_{dz}-1} e^{-(\frac{c_{id}}{\alpha_{dz}})^{\beta_{dz}}}$$

Here, $\hat{\alpha_{dz}}$ and $\hat{\beta_{dz}}$ are the shape and scale parameters that characterize the Weibull distribution. Denoting this distribution $W(\hat{\alpha_{dz}},\hat{\beta_{dz}})$ the estimated distribution $\hat{G_{dz}}$ is formed by combining this with the estimated mass at zero claims, which is the empirical likelihood:

$$\hat{G_{dz}}(c) = \begin{cases} G_{I_{dz}}(0) & \text{if } c = 0 \\ G_{I_{dz}}(0) + \frac{W(\hat{\alpha_{dz}},\hat{\beta_{dz}})(c)}{1-G_{I_{dz}}(0)} & \text{if } c > 0 \end{cases}$$

Again, we use the notation $\hat{G_{iDt}}$ to represent the set of marginal distributions for i over the categories d: the distribution for each d depends on the cell z an individual i is in at t. We combine the distributions $\hat{G_{iDt}}$ for a given i and t into the joint distribution H_{it} using a Gaussian copula method for the mapping Ψ. Intuitively, this amounts to assuming a parametric form for correlation across $\hat{G_{iDt}}$ equivalent to that from a standard normal distribution with correlations equal to empirical rank correlations $\rho_{z_{\theta_{it}}}(D, D')$ described in the previous section. Let $\Phi^i_{1|2|3|4}$ denote the standard multivariate normal distribution with pairwise correlations $\rho_{z_{\theta_{it}}}(D, D')$ for all pairings of the four claims categories D. Then an individual's joint distribution of non-zero claims is:

$$\hat{H_{i,t}}(\cdot) = \Phi_{1|2|3|4}(\Phi_1^{-1}(\hat{G_{id_1t}}), \Phi_2^{-1}(\hat{G_{id_2t}}), \Phi_3^{-1}(\hat{G_{id_3t}}), \Phi_4^{-1}(\hat{G_{id_4t}}))))$$

Above, Φ_d is the standard marginal normal distribution for each d. $\hat{H}_{i,t}$ is the joint distribution of claims across the four claims categories for each individual in each time period. After this is estimated, we determine our final object of interest $F_{jkt}(\cdot)$ by simulating K multivariate draws from $\hat{H}_{i,t}$ for each i and t, and passing these values through the plan-specific total claims to out of pocket mapping Ω_k and the individual to family out of pocket mapping Γ_k. The simulated $F_{jkt}(\cdot)$ for each j, k, and t is then used as an input into estimation of the choice model.

Table B1 presents summary results from the cost model estimation for the final choice model sample, including population statistics on the ACG index θ, the Weibull distribution parameters $\hat{\alpha_{dz}}$ and $\hat{\beta_{dz}}$ for each category d, as well as the across category rank correlations $\rho_{z_{\theta_{it}}}(D, D')$. These are the fundamentals inputs used to generate F_{jkt}, as described above, and lead to accurate characterizations of the overall total cost and out-of-pocket cost distributions (validation exercises which are not presented here).

Final Sample				
Cost Model Output				
	Overall	PPO_{250}	PPO_{500}	PPO_{1200}
Individual Mean (Median)				
Unscaled ACG Predictor				
Mean		1.42	0.74	0.72
Median		0.83	0.37	0.37
Pharmacy: Model Output				
Zero Claim Pr.	0.35 (0.37)	0.31 (0.18)	0.40 (0.37)	0.42 (0.37)
Weibull α	1182 (307)	1490 (462)	718 (307)	596 (307)
Weibull β	0.77 (0.77)	0.77 (0.77)	0.77 (0.77)	0.77 (0.77)
Mental Health				
Zero Claim Pr.	0.88 (0.96)	0.87 (0.96)	0.90 (0.96)	0.90 (0.96)
Weibull α	1422 (1295)	1447 (1295)	1374 (1295)	1398 (1295)
Weibull β	0.98 (0.97)	0.99 (0.97)	0.98 (0.97)	0.98 (0.97)
Hospital / Physician				
Zero Claim Pr.	0.23 (0.23)	0.21 (0.23)	0.26 (0.23)	0.26 (0.23)
Weibull α	2214 (1599)	2523 (1599)	1717 (1599)	1652 (1599)
Weibull β	0.58 (0.55)	0.59 (0.55)	0.55 (0.55)	0.55 (0.55)
($> \$40,000$) Claim Pr.	0.02 (0.01)	0.02 (0.01)	0.01 (0.01)	0.01 (0.01)
Physician OV				
Zero Claim Pr.	0.29 (0.20)	0.26 (0.20)	0.33 (0.46)	0.34 (0.46)
Weibull α	605 (553)	653 (553)	517 (410)	529 (410)
Weibull β	1.15 (1.14)	1.15 (1.14)	1.15 (1.14)	1.14 (1.14)
Correlations				
Rank Correlation Hospital-Pharm.	0.28 (0.34)	0.26 (0.32)	0.31 (0.34)	0.32 (0.34)
Rank Correlation Hospital-OV	0.73 (0.74)	0.72 (0.74)	0.74 (0.74)	0.74 (0.74)
Rank Correlation Pharm.-OV	0.35 (0.41)	0.33 (0.37)	0.38 (0.41)	0.39 (0.41)

Table B1: This table describes the output of the cost model in terms of the means and medians of individual level parameters, classified by the plan actually chosen. These parameters are aggregated for these groups but have more micro-level groupings, which are the primary inputs into our cost projections in the choice model. Weibull α, Weibull β, and Zero Claim Probability correspond to the cell-specific predicted total individual-level health expenses as described in more detail in Appendix B.

C Appendix: Choice Model Estimation Algorithm Details and Additional Results

This appendix describes the details of the choice model estimation algorithm. The corresponding section in the text provided a high-level overview of this algorithm and outlined the estimation assumptions we make regarding choice model fundamentals and their links to observable data. In addition, after the presentation of the estimation algorithm, we discuss further specification details and results for our primary choice model.

We estimate the choice model using a random coefficients simulated maximum likelihood approach similar to that summarized in Train (2009). The simulated maximum likelihood estimation approach has the minimum variance for a consistent and asymptotically normal estimator, while not being too computationally burdensome in our framework. Since we use panel data, the likelihood function at the family level is computed for a *sequence* of choices from t_0 to t_2, since inertia implies that the likelihood of a choice made in the current period depends on the choice made in the previous period. The maximum likelihood estimator selects the parameter values that maximize the similarity between actual choices and choices simulated with the parameters.

First, the estimator simulates Q draws from the distribution of health expenditures output from the cost model, F_{jkt}, for each family, plan, and time period. These draws are used to compute plan expected utility conditional on all other preference parameters. It then simulates S draws for each family from the distributions of the random coefficients γ_j and δ_j, as well as from the distribution of the preference shocks ϵ_k. We define the set of parameters θ as the full set of ex ante model parameters (before the S draws are taken):

$$\theta \equiv (\mu, \beta, \sigma^2_\gamma, \mu_\delta(A_j), \sigma_\delta(A_j), \alpha, \mu_{\epsilon_K}(A_j), \sigma_{\epsilon K}(A_j), \eta_0, \eta_1).$$

We denote θ_{sj} one draw derived from these parameters for each family, including the parameters constant across draws:

$$\theta_{sj} \equiv (\gamma_j, \delta_j, \alpha, \epsilon_{KT}, \eta_0, \eta_1)$$

Denote θ_{Sj} the set of all S simulated draws for family j. For each θ_{sj} the estimator then uses all Q health draws to compute family-plan-time-specific expected utilities U_{sjkt} following the choice model outlined in earlier in section 3. Given these expected utilities for each θ_{sj}, we simulate the probability of choosing plan k in each period using a smoothed accept-reject function with the form:

$$Pr_{sjt}(k = k^*) = \frac{\left(\frac{\frac{1}{-U_{sjk^*t}}(\cdot)}{\Sigma_K \frac{1}{-U_{sjkt}}(\cdot)}\right)^\tau}{\Sigma_k\left(\frac{\frac{1}{-U_{s\hat{j}kt}}(\cdot)}{\Sigma_K \frac{1}{-U_{sjkt}}(\cdot)}\right)^\tau}$$

This smoothed accept-reject methodology follows that outlined in Train (2009) with some slight modifications to account for the expected utility specification. In theory, conditional on θ_{sj}, we would want to pick the k that maximizes U_{jkt} for each family, and then average over S to get final choice probabilities. However, doing this leads to a likelihood function with flat regions, because for small changes in the estimated parameters θ, the discrete choice made does not change. The smoothing function above mimics this process for CARA utility functions: as the smoothing parameter τ becomes large the smoothed Accept-Reject simulator becomes almost identical to the true Accept-Reject simulator just described, where the actual utility-maximizing option is chosen with probability one. By choosing τ to be large, an individual will always choose k^* when $\frac{1}{-U_{jk^*t}} > \frac{1}{-U_{jkt}} \forall k \neq k^*$. The smoothing function is modified from the logit smoothing function in Train (2009) for two reasons (i) CARA utilities are negative, so the choice should correspond to the utility with the lowest absolute value and (ii) the logit form requires exponentiating the expected utility, which in our case is already the sum of exponential functions (from CARA). This double exponentiating leads to computational issues that our specification overcomes, without any true content change since both models approach the true Accept-Reject function.

Denote any sequence of three choices made as k^3 and the set of such sequences as K^3. In the limit as τ grows large the probability of a given k^3 will either approach 1 or 0 for a given simulated draw s and family j. This is because for a given draw the sequence (k_1, k_2, k_3) will either be the sequential utility maximizing sequence or not. This implicitly includes the appropriate level of inertia by conditioning on previous choices within the sequential utility calculation. For example, under θ_{sj} a choice in period two will be made by a family j only if it is optimal conditional on θ_{sj}, other preference factors, and the inertia implied by the period one choice. For all S simulation draws we compute the optimal sequence of choices for k with the smoothed Accept-Reject simulator, denoted k^3_{sj}. For any set of parameter values θ_{Sj} the probability that the model predicts k^3 will be chosen by j is:

$$\hat{P_j^{k^3}}(\theta, F_{jkt}, \mathbf{Z}_j^A, \mathbf{Z}_j^B, H_j, A_j) = \Sigma_{s \in S} \mathbf{1}[k^3 = k^3_{sj}]$$

Let $\hat{P_j^{k^3}}(\theta)$ be shorthand notation for $\hat{P_j^{k^3}}(\theta, F_{jkt}, Z_j^A, Z_j^B, H_k, A_j)$. Conditional on these probabilities for each j, the simulated log-likelihood value for parameters θ is:

$$SLL(\theta) = \Sigma_{j \in J} \Sigma_{k^3 \in K^3} d_{jk^3} \ ln \hat{P_j^{k^3}}$$

Here d_{jk^3} is an indicator function equal to one if the actual sequence of decisions made by family j was k^3. Then the maximum simulated likelihood estimator (MSLE) is the value of θ in the parameter space Θ that maximizes $SLL(\theta)$. In the results presented in the text, we choose $Q = 100$, $S = 50$, and $\tau = 6$, all values large enough such that the estimated parameters vary little in response to changes.

C.1 Specification for Inertia

In the main text we did not describe the details for our specification for consumer inertia. The model for inertia, which is similar to that in Handel (2013), specifics an inertial cost $\eta(\mathbf{Z}_j^B)$ that is linearly related to consumer characteristics and linked choices, $\mathbf{Z}_j^B$:

$$\eta(\mathbf{Z}_j^B) = \eta_0 + \eta_1 Z_{jt}^B$$

The characteristics in $\mathbf{Z}_j^B$ include family status (e.g., single or covering dependents), income, several job status measures, linked choice of Flexible Spending Account (FSA), and whether the family has any members with chronic medical conditions (and, if so, how many chronic conditions total in the family).

C.2 Additional Results

In the interest of space, the text only presented the risk preference parameter estimates from our primary specification, since this was the key object of interest recovered there for our equilibrium analysis of insurance exchange pricing regulations. Here, for completeness, in Tables C1 and C2 we include the full set of estimates in the primary model for reference, including inertia parameters, PPO_{1200} random coefficients, and ε standard deviations. Overall, the parameters not discussed in the text have similar estimates to those in Handel (2013), though the risk preference estimates differ here because they are linked explicitly to health risk to estimate correlations between those two micro-foundations.

Empircal Model Results		
	(1)	Paramater
Parameter / Model	Primary Model	Standard Error
Risk Preference Estimates		
μ_γ - Intercept, β_0	$1.21 * 10^{-3}$	$5.0 * 10^{-5}$
μ_γ - $\log(\Sigma_{i\epsilon j}\lambda_i)$, β_1	$-1.14 * 10^{-4}$	$9.8 * 10^{-6}$
μ_γ - age, β_2	$-5.21 * 10^{-6}$	$1.0 * 10^{-7}$
μ_γ - $\log(\Sigma_{i\epsilon j}\lambda_i)*$age, β_3	$1.10 * 10^{-6}$	$1.3 * 10^{-7}$
μ_γ - Manager, β_4	$4.3 * 10^{-5}$	$5.2 * 10^{-5}$
μ_γ - Manager ability, β_5	$1.4 * 10^{-5}$	$1.2 * 10^{-5}$
μ_γ - Non-manager ability , β_6	$7.5 * 10^{-6}$	$2.4 * 10^{-6}$
μ_γ - Population Mean	$4.39 * 10^{-4}$	-
μ_γ - Population σ	$6.63 * 10^{-5}$	-
σ_γ - γ standard deviation	$1.24 * 10^{-4}$	$3.5 * 10^{-5}$
Inertia Estimates		
η_0, Intercept	1,336	76
η_1, Family	2,101	52
η_1, FSA Enroll	-472	44
η_1, Income	96	15
η_1, Quantitative	6	27
η_1, Manager	162	34
η_1, Chronic Condition	108	24

Table C1: This table presents the first half of the full set of primary choice model estimates: the set of estimates relevant for our analysis of exchange pricing regulation is presented and interpreted in much more detail in the main text. Standard errors are presented in column 2.

Empircal Model Results		
	(1)	Parameter
Parameter / Model	Primary Model	Standard Error
PPO$_{1200}$ Preferences		
μ_δ : Single	-2,504	138
σ_δ : Single	806	47
μ_δ : Family	-2,821	424
σ_δ : Family	872	48
Other		
α, High-Cost, PPO$_{250}$	-805	79
ε_{500} , σ_ε, Single	50	340
ε_{1200} , σ_ε, Single	525	180
ε_{500} , σ_ε, Family	141	56
ε_{1200} , σ_ε, Family	615	216

Table C2: This table presents the second half of the full set of primary choice model estimates: the set of estimates relevant for our analysis of exchange pricing regulation is presented and interpreted in much more detail in the main text. Standard errors are presented in column 2.

D Appendix: MEPS Analysis Descriptives

This section presents some extra tables to support the analysis that re-weights our population according to demographics in the nationally representative MEPS data. See Section 6 in the text for our primary equilibrium and welfare analysis using these re-weighted data. Table D.0 describes the number of individuals in MEPS in each year we use (these years overlap exactly with those from our data). Table D.1 presents detailed characteristics of our population of interests (i) all individuals in MEPS (ii) all individuals in MEPS 25-65 and (iii) all uninsured / individual market insured individuals in MEPS, age 25-65. Table D.2 describes the insurance coverage statistics for each of these three sample. Table D.3 describes the weights used to re-weight our own data for the analysis in the text, while Table D.4 provides a detailed breakdown of health status for these three populations.

Table D.0: MEPS Data Sample Counts

Year	Number of Individuals
2004	34,403
2005	33,961
2006	34,145
2007	30,964
2008	33,066

	Entire MEPS (1)	All Ind. 25-65 (2)	25-65 Unins/Ind (3)
N - Individual-Year Obs.	166,539	81,733	21,856
N - Individuals in Panel	105,353	51,922	13,804
N - Family-Year Obs.	58,647	-	-
N - Families in Panel	36,317	-	-
Avg. Family Members	2.90	-	-
Age-Individual			
Mean	33.82	43.15	42.6
10th Qtile	5	28	27
25th Qtile	14	34	32
Median	32	43	42
75th Qtile	51	52	52
90th Qtile	66	59	60
Gender-Individual			
Male %	47.7%	46.6%	50.2%
Total Income-Family-Year* **			
Mean	53613	64058	42746
10th Qtile	9240	12733	8000
25th Qtile	19000	26000	17068
Median	39080	50000	31114
75th Qtile	72375	85584	54995
90th Qtile	115086	131080	89600
Wage Income-Family-Year**			
Mean	44583	59945	38882
10th Qtile	0	7348	300
25th Qtile	8000	24000	14280
Median	32000	48300	30000
75th Qtile	65000	83753	52000
90th Qtile	104438	124996	82680
Region-Individual			
Northeast	14.5%	15.0%	10.1%
Midwest	19.2%	19.6%	15.0%
South	38.3%	38.7%	46.3%
West	26.9%	26.8%	28.7%

Table D1: This table describes demographic data for key samples of interest in the MEPS data, for the pooled data from 2004-2008. A more detailed description of each column's sample is contained in the text.

*In individual samples, a given family's income may count twice since two individuals can be from same family.

	Entire MEPS (1)	All Ind. 25-65 (2)	25-65 Unins/Ind (3)
Family-Year: Coverage Type*			
Private (Employer or Ind.)	66.3%	73.3%	41.0%
Medicaid (someone)	30.7%	33.4%	45.4%
Medicare (someone)	29.01%	14.0%	16.4%
Uninsured** (someone)	26.7%	35.0%	84.7%
Only Public in Fam	22.5 %	15.1%	0%
Always Offered Employer (someone)	48.8 %	62.1%	–
Offered Employer Sometimes (someone)	62.0%	76.1%	–
Family Member Emp. Always	69.7%	84.7%	76.2%
Family Member Emp. Once	77.5%	92.3%	87.4%
Individual-Year: Coverage Type*			
Private (Employer or Ind.)	54.5%	64.0%	16.8%
Medicaid	25.4%	12.4%	0.72%
Medicare	13.4%	3.9% 1	.25%
Uninsured**	16.6%	22.3%	83.2%
Only Public	27.6%	12.7%	0%
Always Offered Employer	21.3 %	38.9%	–
Offered Employer Sometimes	32.5%	55.0%	–
Individual Emp. Always	37%	65.4%	37.5%
Individual Emp. Once	48%	78.3%	48.0%

Table D2: This table describes insurance coverage, expenditures, and other statistics in the MEPS data for the pooled data from 2004-2008. A more detailed description of each column's sample is contained in the text.

*Coverage type reflects whether a family ever had this kind of coverage (for any member) throughout the year, so these numbers add to more than 100%.

**Uninsured variable occurs when none of other coverage types are held, and the family is uninsured for whole year.

MEPS Weights Incorporated					
All 25-65 Sample					
Age Bucket / Fam. Wages	0-$35,000	$35,000-$70,000	$70,000-$105,000	$\geq$ 105,000	**Total**
25-29	4.1%	4.5	2.7	1.9	13.1%
30-34	3.3%	4.4	2.6	1.9	12.3%
35-39	3.5%	4.2	2.8	2.3	12.9%
40-44	3.6%	4.5	3.0	2.8	13.9%
45-49	3.5%	4.2	3.0	3.1	13.9%
50-54	3.5%	3.8	2.8	2.9	13.1%
55-59	3.8%	3.2	2.3	2.3	11.7%
60-64	4.4%	2.3	1.3	1.2	9.2%
Total	29.7%	31.1%	20.5%	18.4%	100%
% Male by Income*	45.6%	49.9%	50.3%	51.4%	

25-65 Unins./ Private					
Age Bucket / Fam. Wages	0-$35,000	$35,000-$70,000	$70,000-$105,000	$\geq$ 105,000	**Total**
25-29	7.4%	5.0	1.9	1.6	15.9%
30-34	6.0%	4.4	1.3	0.7	12.4%
35-39	6.4%	3.5	1.1	0.6	11.6%
40-44	6.1%	4.0	1.4	0.8	12.2%
45-49	6.2%	3.1	1.6	0.9	10.8%
50-54	5.9%	2.9	1.1	0.9	10.8%
55-59	7.0%	2.5	1.1	0.8	11.4%
60-64	10.1%	2.3	0.8	0.8	14.0%
Total	55.1%	27.7%	10.3%	7.1%	100%
% Male by Income*	51.4%	56.2%	55.4%	56.8%	

Table D3: This table describes the discrete age probabiliities for different age / gender / income categories for (i) all individuals in MEPS, age 25-65, and (ii) all uninsured / individual market insured individuals in MEPS, age 25-65. These weights incoporate MEPS sample weights as well, as an additional weighting factor.

*Percentages of gender across age are essentially constant conditional on income, which is why those figures are not presented here.

MEPS Weights Incl.							
All 25-65 Sample							
Age Bucket / Quantile	10th	25th	50th	75th	90th	95th	**Mean**
25-29	0 (0)	0 (203)	125 (843)	620 (2833)	2109 (7638)	4155 (12007)	**997 (2820)**
30-34	0 (0)	0 (241)	224 (940)	922 (3179)	2815 (9040)	5582 (13122)	**1376 (3146)**
35-39	0 (0)	0 (239)	331 (925)	1314 (2928)	3499 (8158)	6333 (13595)	**1696 (3126)**
40-44	0 (0)	25 (258)	450 (967)	1669 (2955)	4513 (7844)	9099 (13843)	**2235 (3544)**
45-49	0 (0)	115 (365)	703 (1342)	2425 (3827)	6423 (9143)	12125 (15505)	**3016 (3838)**
50-54	0 (90)	221 (563)	1114 (1860)	3385 (4744)	8562 (10683)	16271 (17135)	**4187 (4551)**
55-59	0 (102)	410 (781)	1837 (2437)	4953 (5820)	11929 (13615)	21069 (22741)	**5315 (6129)**
60-64	71 (255)	707 (1109)	2337 (2906)	5916 (6771)	15261 (14493)	27033 (24997)	**6790 (6666)**

25-65 Unins./ Private							
Age Bucket / Quantile	10th	25th	50th	75th	90th	95th	**Mean**
25-29	0 (0)	0 (0)	0 (166)	173 (758)	819 (2959)	1824 (5502)	**391 (952)**
30-34	0 (0)	0 (0)	0 (180)	254 (852)	1062 (3234)	2024 (6095)	**608 (1322)**
35-39	0 (0)	0 (0)	0 (174)	328 (1024)	1650 (3187)	3164 (5748)	**744 (1223)**
40-44	0 (0)	0 (0)	50 (308)	750 (1459)	2929 (3966)	4500 (6908)	**1381 (2449)**
45-49	0 (0)	0 (0)	120 (425)	857 (1846)	3108 (4566)	6719 (9658)	**2089 (1967)**
50-54	0 (0)	0 (144)	340 (798)	1576 (2866)	5590 (7462)	11851 (12952)	**2474 (3085)**
55-59	0 (0)	24 (176)	1076 (1312)	3565 (3996)	9290 (9990)	16419 (19459)	**3898 (4941)**
60-64	0 (60)	449 (732)	1966 (2398)	5166 (5730)	13749 (12017)	24157 (21839)	**6003 (6043)**

Table D4: This table describes the expenditure quantiles for (i) all individuals in MEPS age 25-65 (top panel) and (iii) all uninsured / individual market insured individuals in MEPS, age 25-65 (bottom panel). Female numbers presented in parantheses, male numbers are not.