Imperfect Risk Adjustment, Risk Preferences, and Sorting in Competitive Health Insurance Markets*

Timothy Layton
Boston University

February 27, 2014

Abstract

Much of the risk adjustment literature focuses on its effects on insurers' incentives to inefficiently manipulate insurance contracts to "cream-skim" the healthiest enrollees in the market (Glazer and McGuire 2000). However, when prices are set competitively as in the Exchanges established by the ACA, risk adjustment can also ameliorate another important type of selection problem, in the extreme case known as market unraveling or death spirals, where consumers inefficiently sort between plans due to the correlation between costs and demand. In this paper, I study how imperfect risk adjustment affects prices, sorting, and welfare in competitive health insurance markets. First, I build on the model of Einav et al. (2010) to show graphically and in a theoretical model that imperfect risk adjustment causes plan prices to be based on the portion of costs not predicted by the risk adjustment model ("residual costs"), rather than total costs. This adjusting of costs has the potential to substantially weaken the connection between demand and costs that is the source of the adverse selection problem, and the extent to which it will do so depends on the correlation between demand and the costs predicted by the risk adjustment model ("predicted costs"). In a setting where consumers are required to choose between two plans, and one plan is adversely selected, if predicted costs are positively (negatively) correlated with demand for that plan, risk adjustment will cause the prices of the two plans to converge (diverge), resulting in more (fewer) consumers choosing the adversely selected plan. I then use administrative health insurance claims data from a large employer to estimate the joint distribution of preferences, total costs, and predicted costs for a large sample of employees. I use the estimates to simulate competitive equilibria under several common forms of risk adjustment in a setting similar to the Exchanges. I find that in this setting when there is no risk adjustment, the market completely unravels and the entire market enrolls in the less comprehensive plan. However, when diagnostic-based risk adjustment similar to that being implemented in the Exchanges is implemented, a substantial portion of the market unraveling is undone, with over 80% of individuals enrolling in the more comprehensive plan. Estimates suggest that this implies a substantial welfare gain of over $700 per person, per year in this setting.

1 Introduction

The question of whether competition improves efficiency in health insurance markets is at the center of the recent reform of the US health care system. Efficiency is achieved when consumers purchase goods that they value more than

* **Acknowledgments:** I am indebted to my advisers, Randy Ellis, Tom McGuire, and Keith Ericson, for their guidance and support in the preparation of this paper. I thank Claudia Olivetti helping me obtain access to the data for this project at the National Bureau for Economic Research. for I am also grateful to Francesco Decarolis, Daria Pelech, Mark Sheppard, and Julie Shi for useful comments and suggestions. I thank Mike Geruso and Ben Handel for incredibly helpful methodological advice. I also thank seminar participants at the Department of Health Care Policy at Harvard Medical School for comments and suggestions. I gratefully acknowledge financial support from the Institute for Economic Development at Boston University and the National Institute of Mental Health (R01 MH094290).

the cost of those goods. In most markets, competition induces efficiency by ensuring that goods are priced according to their marginal cost. In many health insurance markets, however, competition does not ensure that the price an individual pays for a product (an insurance plan) is equal to the marginal cost of providing that product. This is due to the fact that in insurance markets, the cost of the product depends on the characteristics of the consumer purchasing it, and asymmetric information or regulation often precludes an insurer from matching the price an individual pays to her expected cost.[1] In a competitive equilibrium, the price a consumer pays for an insurance policy will reflect the average cost of all consumers choosing to purchase the policy. This price may be quite different from the marginal cost to the insurer of enrolling the consumer. When the premium paid by a consumer differs from her expected cost, consumers may sort inefficiently across plans (Akerlof 1970, Einav et al. 2010). In the extreme case, some insurance plans may cease to exist due to market unraveling or death spirals (Cutler and Reber 1998). This market failure is known as adverse selection, and due to competitive pressures and a less-restricted contract space, the Health Insurance Exchanges (Exchanges) established by the Affordable Care Act (ACA) are likely to experience much more serious adverse selection problems than in other settings such as the employer and Medicare markets more typically studied in the literature. In fact, some recent research suggests the potential for complete market unraveling in an Exchange-like setting (Handel et al. 2013).

The potential for adverse selection in insurance markets is widely recognized among economists, and has been widely studied empirically.[2] However, the health economics literature has largely focused on only three solutions: restricting the contract space, subsidizing adversely selected plans, and allowing premiums to vary to some extent by expected cost (Cutler and Reber 1998, Einav et al. 2010, Bundorf et al. 2012, Geruso 2013, Handel et al. 2013). Risk adjustment, one of the most prevalent tools for combatting adverse selection, has been mostly overlooked.[3] This has largely been due to the fact that risk adjustment is typically used to combat supply-side selection problems, known as cream-skimming, where insurers manipulate insurance contracts to attract healthy enrollees.[4] In this paper, I show that risk adjustment will also have significant effects on efficiency in a fixed contract setting where plans cannot manipulate their contracts and inefficiency is due to sorting rather than cream-skimming. The choice of plan tier in the Exchanges represents such a setting.

[1] In fact, in the US the Affordable Care Act (ACA) prohibits virtually all variation in health insurance premiums across consumers, forcing insurers to charge one price to all consumers, no matter their expected costs.

[2] see Einav et al. 2011 for a recent review

[3] A few recent papers have also studied the effects of risk adjustment on demand-side adverse selection. Glazer et al. (2013) develop a risk adjustment model that induces optimal sorting when there is adverse selection using constrained regression. Shi (2013) studies the interaction of age-based pricing and risk adjustment in the context of a simulated Exchange. Handel et al. (2013) simulate a competitive equilibrium under a from of perfect risk adjustment based on ex-post realized costs rather than the risk scores used in practice.

[4] see van de Ven and Ellis 2000 and Ellis and Layton 2013 for reviews of this literature

Risk adjustment is a policy that transfers costs across plans according to the health status of each plan's enrollees. It has been implemented in some form in almost every individual health insurance market in the world, including the new state Exchanges.[5]It works by first calculating a risk score for each individual in the market. The risk score is a proxy for the individual's total expected medical spending. If a more comprehensive plan attracts enrollees with high risk scores while a less comprehensive plan attracts enrollees with low risk scores, risk adjustment will result in a transfer from the less comprehensive plan to the more comprehensive plan. It is straightforward to see how such a policy might affect plan prices and sorting. Consider the case of "perfect" risk adjustment where risk scores perfectly predict each individual's total expected medical spending. In this case, the transfers will eliminate all individual-level cost variation. This will result in all enrollees having identical total costs from the plan's point of view. Plan average costs will then depend only on the portion of total cost covered by the plan, not on the health status of the plan's enrollees. If plans set prices equal to average cost, as they would in a competitive equilibrium, this shift in average cost will potentially have significant effects on prices and sorting of consumers across plans. Considering the large differences in risk protection offered by the most and least comprehensive plans in the Exchanges (90% vs. 60% of health care costs for a typical enrollee), this re-sorting of consumers across plans could potentially have significant consequences for social welfare.

In practice, however, risk adjustment is imperfect. Risk scores are constructed from somewhat "exogenous" predictors of health status such as demographics and diagnoses. These predictors explain only a small portion of the variance in individual costs. While perfect risk adjustment would effectively pool *all* costs among the plans in a market, imperfect risk adjustment only pools the specific dimensions of costs correlated with the predictors used to construct the risk scores. This is important because medical spending is a multi-dimensional object. Consider a setting where individuals choose between two health plans varying only in cost sharing, Plan E (enhanced) and Plan B (basic). Some dimensions of spending, such as spending on chronic diseases, are likely to be highly correlated with demand for Plan E. This will result in higher spending on chronic diseases in the more comprehensive plan. Other dimensions, such as spending on emergency care for accident victims, may be negatively correlated with demand for Plan E if individuals who are more accident-prone tend to be less risk averse (Cutler et al. 2008). Suppose spending on chronic diseases is much greater than spending on emergency care so that total spending is higher in Plan E. In competition, price will be equal to average cost, so Plan E will have a higher price. Now suppose risk scores are based on indicators for chronic diseases. In this case, when comparing the risk scores of the enrollees of plans E and B, Plan E will appear to have sicker enrollees. Risk adjustment will result in transfers from Plan B, the less comprehensive plan, to Plan E, the more comprehensive plan. Plan E's cost will be reduced and Plan B's cost will be increased by the transfer. This will result in a convergence of the premiums for Plan E and Plan B. Now, suppose risk scores are based on indicators

[5]In the US, risk adjustment is used in some form in Medicare Advantage, Medicare Part D, the new state Health Insurance Exchanges, and many state Medicaid Managed Care programs. Risk adjustment is also used in some form in the health insurance markets of the Netherlands, Switzerland, Germany, Israel, and Belgium.

for accidents or emergency room (ER) use. In this case, when comparing the risk scores of the enrollees of plans E and B, Plan B will appear to have sicker enrollees. Now, risk adjustment will counterintuitively result in transfers from Plan E, the *more* comprehensive plan, to Plan B, the *less* comprehensive plan and a divergence of prices. In both cases, preferences and demand are the same, but the effect of risk adjustment on prices will be different. The different outcome is due to the choice of predictors used to construct the risk score.

In the first part of this paper, I present a novel framework for evaluating the effect of imperfect risk adjustment on prices and sorting in competitive health insurance markets. I develop a model, based on the model of Einav et al. (2010). The key innovation of the model is to divide individual health care costs into two dimensions: predicted costs and residual costs. Predicted costs are the portion of individual costs that are correlated with the predictors used to construct the risk scores. Residual costs are all other costs. After developing the model, I use a series of graphical representations, building on those presented in Einav et al. (2010) and Einav and Finkelstein (2011), to develop intuition for how risk adjustment affects sorting and welfare. Both the model and the graphical representations are useful in that they show that the effect of risk adjustment on prices and sorting in a particular market can be summarized by two simple statistics: The correlation between demand and predicted cosst and the correlation between demand and total costs. These correlations can be easily recovered given exogenous variation in prices or a structural model of health plan choice.

In the second part of this paper, I investigate the efficiency consequences of risk adjustment empirically using administrative health insurance claims data from a large employer. I use the data to estimate the joint distribution of total costs, predicted costs, and risk preferences for a sample of employees. Following the implications of the model, I allow for preferences to be correlated with both predicted and total costs. I then use this joint distribution to simulate plan prices and consumer sorting under various forms of risk adjustment in the context of a Health Insurance Exchange where prices are set competitively and all consumers choose between a Bronze plan and a Platinum plan.[6] I use these simulations to compare sorting and welfare with no risk adjustment and under a number of risk adjustment models that are being used or could potentially be used in the Exchanges. I simulate risk adjustment based on demographic predictors and based on a variety of commonly used sets of diagnostic-based predictors.

The simulations suggest that there will be significant adverse selection in the Exchanges. I replicate the result of Handel et al. (2013) that with no risk adjustment, the market fully unravels, and all consumers enroll in the less comprehensive Bronze plan. Not surprisingly, I find that in the Exchange setting under all forms of risk adjustment predicted costs are positively correlated with demand for the more comprehensive Platinum plan. This implies that risk adjustment should cause the premiums of the Bronze and Platinum plans to converge. However, under demographic-based risk adjustment, the correlation between predicted costs and demand is not strong enough to undo market

[6] In the Exchanges, plans are divided into tiers based on their actuarial value. The tiers are called (from least to most comprehensive) Bronze, Silver, Gold, and Platinum.

unravelling. On the other hand, diagnostic-based risk adjustment, similar to that being implemented in the Exchanges, eliminates much of the correlation between total costs and demand and *almost fully undoes market unraveling*. The equilibrium premiums of the two plans converge and over 80% of market participants enroll in the more comprehensive Platinum plan. Moreover, when risk adjustment is combined with reinsurance, as it is in the Exchanges, virtually all relevant variation in plan costs across consumers is eliminated and close to 100% of the market enrolls in the Platinum plan. Welfare calculations indicate that the welfare consequences of risk adjustment in this setting are far from trivial, with risk adjustment improving welfare by over $700 per person, per year, or around 20% of total health care costs among employees of the firm I study.[7]These findings are largely robust to various assumptions about the ability of consumers to predict their future costs.

These findings represent an important contribution to the literature on adverse selection in markets for health insurance. First, the effect of imperfect risk adjustment on prices and sorting is theoretically ambiguous, and the direction of the effect depends on the correlation between demand and total costs and the correlation between demand and predicted costs. Second, risk adjustment, a policy meant to limit incentives for plans to cream-skim healthy enrollees, also has a large and important effect on equilibrium prices and sorting in competitive health insurance markets. In fact, in the setting studied here, it proves critical for the market to function efficiently. Additionally, when combined with reinsurance, market unraveling is completely undone. This suggests that risk adjustment may play a much more important role in these markets than was previously assumed and should be taken much more seriously among economists as a solution not just to cream-skimming problems, but also inefficient sorting. These findings also suggest that because risk adjustment can almost completely undo market unraveling, it is unwise to ignore it in any empirical study of the markets where it is being used, such as Medicare Part D and the Exchanges, as it is likely to have large and important effects on plan pricing. At the same time, these results suggest that under proposed risk adjustment policies, residual costs remain important, implying that it is also unwise to assume that risk adjustment is perfect in these markets and that there is no remaining relevant variation in costs.

The paper proceeds as follows. Section 2 develops a simple model of a competitive health insurance market with risk adjustment and presents the graphical framework to provide intuition for the relationship between risk adjustment, prices, and sorting, focusing on the importance of the correlation between demand and predicted costs and the correlation between demand and residual costs. Section 3 discusses the data used for estimation. Section 4 outlines the structural empirical model used to the joint distribution of demand, total costs, and predicted costs. Section 5 presents the results of the simulations of equilibrium under risk adjustment, and Section 6 concludes.

[7]This is a huge welfare improvement. It is worth noting that it is especially large compared to the calculations of welfare loss from adverse selection found elsewhere in the literature (Cutler and Reber 1998, Einav et al. 2010, Geruso 2013). It is important to note, however, that in all of these other settings, the plans consumers were choosing from were quite similar in terms of cost sharing. Here, the plans have huge differences in cost sharing, reflecting the huge differences in cost sharing found across tiers in the Exchanges. Simulations with plan options that are more similar to the options available in the settings studied in other papers found welfare results similar to the results from those papers, suggesting that if the estimated structural demand and cost parameters from those papers were used to study the Bronze-Platinum setting studied here, tthe results would be similar.

2 Theoretical Framework

The model I develop in this paper builds on those developed in Einav et al. (2010) and Bundorf et al. (2012). The key innovation is that total costs are divided into two components: predicted costs and residual costs. In the model individuals are required to choose one of two insurance contracts. Everyone faces the same price for each contract. One of the contracts provides enhanced coverage (contract E) and the other provides basic coverage (contract B). As in Bundorf et al. (2012), consumers are distinguished by their health risk, θ, and preferences, ε. Let $v^j(\theta_i, \varepsilon_i)$ represent consumer i's valuation of plan j in dollars, so that $\Delta v(\theta_i, \varepsilon_i) = v^E(\theta_i, \varepsilon_i) - v^B(\theta_i, \varepsilon_i)$ represents consumer i's willingness-to-pay for Plan E over Plan B. Let P represent the difference in the price for Plan E and the price for Plan B. Therefore, individual i chooses to purchase Plan E if and only if $\Delta v(\theta_i, \varepsilon_i) \geq P$.

Now, to see how plans set P, let $c^j(\theta_i)$ represent the expected monetary cost to plan j of enrolling an individual with health risk θ_i. Assume that insurers set prices in perfect competition. While this is not likely the case in the Exchanges where I am simulating competitive equilibria, it is a convenient benchmark. In competition, insurers will set prices equal to average cost. Therefore, I describe the average costs of plans E and B as follows:

$$AC^E(P) = E\left[c^E(\theta) \mid \Delta v(\theta, \varepsilon) \geq P\right]$$

$$AC^B(P) = E\left[c^B(\theta) \mid \Delta v(\theta, \varepsilon) < P\right]$$

In equilibrium, the premium differential P will then be equal to the difference between E and B's average costs:[8]

$$P^* = \Delta AC(P) = AC^E(P) - AC^B(P)$$

The expressions show that two factors cause the premium differential, P^*, to vary. To see the first factor, let us assume that the individuals enrolling in Plan E and Plan B are identical. Now, if the cost to Plan E from enrolling individual i is different from the cost to Plan B from enrolling the same individual, i.e. $c^E(\theta_i) \neq c^B(\theta_i)$, P^* will reflect that cost difference. Plan costs for enrolling the same individual could differ for a variety of reasons including differences in plan generosity, moral hazard, administrative costs, etc. To see the second factor, let us now assume the cost to Plan E from enrolling individual i is identical to the cost to Plan B from enrolling the same individual, i.e. $c^E(\theta_i) = c^B(\theta_i)$. Now, if the individuals enrolling in Plan E and Plan B are identical, the average cost for each plan will also be identical, and $P^* = 0$. However, if the individuals enrolling in Plan E are sicker (and thus higher cost) than

[8]It is possible that a set of premiums resulting in positive enrollment in plans E and B will not exist. In this case, as shown in Handel et al. (2013) in equilibrium all individuals will enroll in Plan B if $\Delta AC(P)$ is always greater than θ for all values of θ and in Plan E if $\Delta AC(P)$ is always less than θ for all values of θ in the population. In these cases P^* will be equal to $AC^B(P)$ and $AC^E(P)$, respectively.

the individuals enrolling in Plan B, Plan E's average cost will be larger than Plan B's and the premium differential will be positive, $P^* > 0$.

This situation is illustrated in Figure 1. In the figure, I assume that $c^E(\theta_i) = c^B(\theta_i)\ \forall \theta_i$ and that higher cost individuals prefer Plan E more than lower cost individuals, i.e. θ_i is positively correlated with demand for E. In both panels of Figure 1, enrollment in E is increasing along the x-axis. The left panel shows how AC^E and AC^B change as enrollment in E increases. Both fall because the positive correlation between demand for E and θ implies that each additional individual enrolling in E is healthier than the other individuals in E and sicker than the other individuals in B. Recall, however that it is the difference between these two lines, ΔAC, that matters for the equilibrium price differential, P^*. The relationship between enrollment in E and ΔAC can be seen in the right panel of Figure 1. In the figure, $\Delta AC(P)$, or the incremental average cost curve, is defined by the difference between the cost of the first individual to enroll in E and the average cost of everyone else, ΔAC^1, on the left and the difference between the cost of the last individual to enroll in E and the average cost of everyone else, ΔAC^N, on the right. The equilibrium in this setting occurs where the demand curve crosses $\Delta AC(P)$. This is where the sorting of individuals across plans induced by P results in $P^* = \Delta AC(P^*)$, the competitive equilibrium.

Note that the left panel of Figure 1 suggests that both the slope and the intercept of $\Delta AC(P)$ are affected by the relationship between θ and demand for E. If AC^E and AC^B were downward sloping (implying adverse selection) and parallel, $\Delta AC(P)$ would be flat and the intercept would be equal to the constant difference between the two lines.[9] This suggests that the extent of adverse selection is illustrated by both the slope and the right-hand side intercept of the incremental average cost curve. If the curve has either a smaller intercept or a flatter slope, the point at which it crosses the demand curve will be lower and further to the right, implying more individuals will enroll in Plan E and the price differential will be lower. This is different from the figures in Einav et al. (2010) and Einav and Finkelstein (2011) where consumers choose between insurance and uninsurance and the right-hand side intercept of the incremental average cost curve is fixed, leaving only the slope to describe the extent of adverse selection.

2.1 Risk Adjustment

I now augment the model by adding risk adjustment. To incorporate risk adjustment into the model, I introduce the concept of predicted costs. As explained in the introduction, risk adjustment consists of a set of transfers across health plans in a market. These transfers are based on a proxy for individual health risk, r_i, typically known as the risk score. Effectively, the transfers result in the portion of costs that are correlated with r_i being pooled among all plans in the market. This causes the new incremental average cost under risk adjustment to be based only on the portion of costs not explained by r_i. Formally, let predicted costs be the costs pooled by a particular risk adjustment

[9] Generally, the slope and intercept of $\Delta AC(P)$ are functions of the distribution of the joint distribution of health risk, θ_i, and preferences, ε_i, in the population. If health risk is distributed uniformly and preferences are constant across individuals, then AC^E and AC^B will be parallel but downward sloping. Other distributions will result in non-parallel average cost curves and different intercepts.

policy, $c_P(r_i) = f(r_i)$. . In general, the regulator implements risk adjustment by paying plan j $c_P(r_i)$ and charging each plan $\bar{c}_P = \frac{1}{N}\sum_{i=1}^{N} c_P(r_i)$ for each individual i enrolled in plan j, where N is the number of individuals in the market.[10]Typically, r_i is calculated by multiplying a vector of predictors, X_i, known as "risk adjusters" by a vector of "risk adjustment weights", β, normalized by the average risk score in the population, $\bar{R}_i$:

$$r_i = \frac{X_i\beta}{\bar{R}_i}.$$

For demographic-based risk adjustment, X_i consists of a series of age-by-sex cells. For diagnostic-based risk adjustment, X_i consists of age-by-sex cells and indicators for a set of clinical condition categories. Each indicator is equal to 1 if a diagnosis mapped to the condition category appears in the individual's health insurance claims from either the current (concurrent) or the prior year (prospective).[11] Predicted costs are a linear function of this risk score and the average cost in the population: $c_p(r_i) = r_i \cdot AC$.[12] Therefore, with risk adjustment, plan average costs are

$$\begin{aligned} AC^j_{RA}(P) &= E\left[c^j(\theta_i) - c_P(r_i) + \bar{c}_P \mid \Delta v(\theta,\varepsilon) \geq P\right] \\ &= AC^j(P) - E\left[c_P(r_i) \mid \Delta v(\theta,\varepsilon) \geq P\right] + \bar{c}_P \end{aligned}$$

and the equilibrium price differential is

$$\begin{aligned} P^*_{RA} &= \Delta AC_{RA}(P) \\ &= \Delta AC(P) - \left(E\left[c_P(r_i) \mid \Delta v(\theta,\varepsilon) \geq P\right] - E\left[c_P(r_i) \mid \Delta v(\theta,\varepsilon) < P\right]\right) \\ &= \Delta AC(P) - \left(AC^E_P - AC^B_P\right) \\ &= \Delta AC(P) - \Delta AC_P(P) \end{aligned}$$

As before, P^*_{RA} reflects differences in the cost to Plan E from enrolling individual i and the cost to Plan B from enrolling the same individual. However, whereas before P^* also reflected the relationship between demand for Plan E and total individual health care costs, it now reflects only the relationship between demand and the portion of costs not explained by $c_P(r_i)$, or the "residual" costs. For demographic-based risk adjustment, the residual costs are all costs not predictable using age-by-sex cells. For diagnostic-based risk adjustment, the residual costs are all costs not predictable using age-by-sex cells and diagnosis groups. The general concept is illustrated in Figure 2. Again, enrollment in Plan E

[10]Here, I assume that an individual's predicted costs do not depend on the plan the individual enrolls in. In many settings this is true. However, in the Exchanges, while individuals' risk scores are the same in different plans in the same metal tier, they vary across tiers. I account for this in the empirical section but abstract from it here for simplicity. This assumption would also break down if there is "upcoding" (Geruso and Layton 2014).

[11]A detailed explanation of the demographic- and diagnostic-based risk adjustment models, including the method of calculating the weights, can be found in the appendix.

[12]In the Exchanges, risk adjustment is implemented using plan, rather than individual, transfers. However, a simplified version of the plan transfers are equivalent to individual transfers equal to $t_i = [r_i - 1]\bar{P}$, where $\bar{P}$ is the average premium in the state. Assuming perfect competition, i.e. $\bar{P} = AC$, this is equivalent to the formulation of risk adjustment presented here.

is increasing along the x-axis. While Figure 1 illustrated the relationship between enrollment in E and the average cost in each plan, Figure 2 illustrates the relationship between enrollment in E and the average *predicted* cost in each plan. In the case shown in the figure, average predicted cost is increasing in demand for E. Similar to Figure 1, this implies that there is adverse selection in predicted costs. Note that the extent and direction of selection in predicted costs is entirely dependent on the relationship between the proxy for health risk, r_i, and demand rather than the correlation between actual health risk, θ_i, and demand.

Figure 3 combines Figures 1 and 2 by plotting $\Delta AC(P)$ and $\Delta AC_P(P)$ on the same plot. The figure clearly shows how incremental average costs can be split into incremental predicted and incremental residual costs. Recall that when risk adjustment is implemented, a plan receives a payment equal to the predicted costs of its enrollees and pays the regulator the average predicted cost in the market. This implies that under risk adjustment the difference in predicted costs in Plan E and Plan B, $\Delta AC_P(P)$ will be zero, leaving only differences in residual costs, $\Delta AC_{RA}(P)$. This is illustrated in Figure 4. This figure shows the effect of risk adjustment on equilibrium prices and sorting. Without risk adjustment, the equilibrium occurs at point A where the demand curve crosses $\Delta AC(P)$. With risk adjustment, however, the equilibrium occurs at point B where the demand curve crosses $\Delta AC_{RA}(P)$. The incremental average risk adjusted cost curve, $\Delta AC_{RA}(P)$, is defined on the left by the portion of the difference between the first enrollee's costs and the costs of everyone else not explained by $c_P(r_i)$ and on the right by the portion of the difference between the last enrollee's costs and the costs of everyone else not explained by $c_P(r_i)$. It is $\Delta AC_P(P)$ that determines the effects of risk adjustment on the equilibrium. In the case illustrated in the figure, a about half of the incremental average total costs are explained by $c_P(r_i)$, resulting in $\Delta AC_{RA}(P)$ being flatter and having smaller left-hand and right-hand side intercepts than $\Delta AC(P)$.

In the case presented in Figures 1-4, Plan E is adversely selected on both total and predicted costs. However, it is not necessary that selection on total and predicted costs be in the same direction. As discussed in Section 1, medical spending is multidimensional. Consider a case where total costs are made up of two dimensions: chronic disease costs and emergency costs. Chronic disease costs are likely to be positively correlted with demand due to their predictability. Emergency costs, on the other hand, may be negatively correlated with demand due to unpredictability and a negative correlation with risk preferences. For example, Cutler et al. (2008) showed that risky behaviors that can result in costly emergency care, such as not wearing a seatbelt, tend to be negatively correlated with demand for insurance. If chronic disease costs are larger than emergency costs, then demand will be positively correlated with total cost. If risk scores are based on indicators for chronic diseases, then both predicted costs and total costs will be positively correlated with demand, and Figures 1-4 show the effects of risk adjustment on prices and sorting. However, if risk scores are based on indicators for emergency room use, then total costs will still be positively correlated with demand, but predicted costs will now be negatively correlated with demand. Figure A4 in the appendix illustrates this type of situation graphically. The left panel of the figure shows the average predicted costs in Plan E and Plan B as a function of enrollment in Plan

E. Both are *increasing*, rather than decreasing, in Plan E's enrollment, implying Plan E is advantageously selected on predicted costs. This generates the incremental average predicted cost curve in the right panel of Figure A4. Note that the curve is below the x-axis and upward sloping due to the advantageous selection on predicted costs. This implies that if risk adjustment is implemented, transfers to E will be negative and transfers to B will be positive. This can be seen more clearly in Figure A5. This figure combines Figure A4 with Figure 1. In this case incremental total costs are still divided into two components, incremental predicted and incremental residual costs, but now because incremental predicted costs are negative, incremental residual costs must exceed incremental total costs. Figure A6 shows what happens when risk adjustment is implemented in this situation. Because incremental predicted costs are negative and upward sloping, the new incremental average risk adjusted cost curve is both shifted up and given a steeper slope. In the figure, the shift is extreme enough that the incremental average risk adjusted cost curve is everywhere above the demand curve, implying that with risk adjustment the market will unravel and everyone will enroll in Plan B. Therefore, in this situation risk adjustment has the opposite effect on prices and sorting than in the situation described in Figures 1-4. Here, risk adjustment reduces enrollment in E and increases the incremental price, while in Figures 1-4, risk adjustment increases enrollment in E and decreases the incremental price.

This suggests that both the direction and the size of the effect of risk adjustment on prices and sorting are theoretically ambiguous and therefore an empirical questions. The Figures also show that when the slopes of the average total cost curves in the left panel of Figure 1 are negative, the direction of the effect of risk adjustment on prices and sorting can be determined by observing the sign of the slopes of the average predicted cost curves in the left panel of Figures 2 and A4. Negative slopes (average predicted costs decreasing in Plan E enrollment) imply that risk adjustment will cause prices to converge and enrollment in Plan E to increase. Positive slopes (predicted costs increasing in Plan E enrollment) imply that risk adjustment will cause prices to diverge and enrollment in Plan E to decrease.[13] These slopes can be easily recovered using exogenous variation in prices. They can also be determined using a structural model of health insurance demand.

I note here that when consumers are choosing between two health insurance plans that vary only in cost-sharing (such as choice of tier in the Exchanges) and predicted costs consist of the portion of costs that is predictable using diagnoses, this slope will usually be positive. This is due to the fact that consumers' demand for insurance will largely be based on expectations of their future spending. These expectations will likely be somewhat similar to expected spending based on the diagnoses. This will result in demand being positively correlated with both total and predicted costs. However, in a setting where health plans are more horizontally differentiated by network or quality rather than by cost-sharing, the correlation between different components of cost and demand places a much higher weight on preferences than predictability. This is similar to the point made in Bundorf et al. (2012). This makes it clear that when investigating the equilibrium consequences of risk adjustment empirically, or when studying insurer behavior

[13] When the slopes of the average total cost curves are positive (advantageous selection), these results should be reversed.

in markets where risk adjustment has been implemented, it is important to estimate not only the relationship between preferences and total costs (Finkelstein and McGarry 2006) but also the relationship between preferences and $f(r_i)$.

To tie this framework to the previous literature on risk adjustment, I also note here that $\Delta AC(P)$ characterizes the relationship between demand and total costs while $\Delta AC_P(P)$ characterizes the relationship between demand and predicted costs. In the risk adjustment literature, there is a focus on maximizing the "fit" of a risk adjustment model, or the relationship between r_i and $c(\theta_i)$. However, the model and the figures presented here suggest that in order to reduce or eliminate demand-side selection, it is more important for the relationship between predicted costs and demand to fully characterize the relationship between total costs and demand. While constructing a risk adjustment model that perfectly predicts total costs, i.e. $c_p(r_i) = c(\theta_i)$, would produce this result, other models with lower fit could also achieve the same result. For example, suppose that individuals' demand for E is a linear function of demographic variables, $\Delta v(\theta_i, \varepsilon_i) = \delta r_i^{demo}$. Then, with demographic-based risk adjustment $\Delta AC(P) = \Delta AC_P(P)$. In this case, implementing risk adjustment would eliminate all selection, $\Delta AC_{RA}(P) = 0$. This is true despite the fact that demographic-based risk adjustment models only explain around 1% of the variation in costs in the population. This is because in this case the residual costs, the 99% of the variation in costs not predicted by demographic variables, are not related to demand, and thus are distributed randomly across plans.

2.2 Efficiency

Efficiency requires that an individual enroll in Plan E if and only if her willingness-to-pay for Plan E over Plan B exceeds the incremental marginal cost of enrolling her: $\Delta v(\theta_i, \varepsilon_i) \geq \Delta c(\theta_i) = c^E(\theta_i) - c^B(\theta_i)$. Note that risk adjustment does not affect the efficiency criteria. This is because the social cost of enrolling individual i in plan j is invariant to any transfers of costs across plans: Individual i will always cost the market more in Plan E than in Plan B. Risk adjustment just changes how those costs are distributed across plans in the market.

Recall that individuals sort across plans according to their willingness-to-pay, $\Delta v(\theta, \varepsilon)$. At the same time, efficiency requires them to sort according to their incremental marginal cost, $\Delta c(\theta_i)$. It is important to note that in this environment whether there is risk adjustment or not, there is only one tool to induce sorting, and thus affect welfare: the uniform price differential, P. Risk adjustment affects efficiency by altering the equilibrium value of P, thus causing market participants to re-sort between plans. As noted above, the extent and direction of this re-sorting depends on the correlation between demand and predicted costs. The welfare consequences of this re-sorting depends on the joint distribution of $\Delta v(\theta_i, \varepsilon_i)$ and $\Delta c(\theta_i)$. [14]

The case where $\Delta c(\theta_i)$ and demand vary linearly with P is shown in Figure 5, where the relationship between

[14] In addition to making it unclear how risk adjustment will affect sorting, the potential for preference heterogeneity also makes it unclear ex-ante what optimal sorting looks like (Bundorf et al. 2012). However, a treatment of this issue is beyond the scope of the paper. Instead, I just suggest that the effect of risk adjustment on welfare is unclear ex-ante. A thorough study of risk adjustment in such an environment is a promising and interesting direction for future research.

$\Delta c(\theta_i)$ is represented by the green incremental marginal cost curve, $IMC(P)$. In this case, the welfare loss from adverse selection with no risk adjustment is described by the combination of the two shaded areas. With risk adjustment, the remaining welfare loss from adverse selection is described only by the solid shaded area, implying a large welfare improvement from risk adjustment. When the correlation between predicted costs and demand is negative, however, risk adjustment results in a steepening of $\Delta AC(P)$ which would increase the size of the shaded area describing welfare loss rather than decrease it.

In the figure, and in the empirical case studied below, the incremental marginal cost curve is everywhere below the demand curve, implying that it is optimal for the entire market to enroll in the Platinum plan. In reality, however, a variety of factors such as moral hazard or preference heterogeneity can result in the incremental marginal cost curve crossing the demand curve. When this occurs, the welfare consequences of risk adjustment are even more ambiguous, as risk adjustment could result in "too many" individuals enrolling in the more comprehensive plan.

3 Data and Setting

In order to study the effects of risk adjustment on equilibrium sorting and prices, I estimate the joint distribution of preferences, total costs, and predicted costs using data from a large employer in the Truven Marketscan Database during 2006-07.[15] During this period, the firm offered its employees a choice of two PPO plans: a basic plan (Plan B) and a more comprehensive, enhanced plan (Plan E). Around 50,000 employees enrolled in these plans during the time period, along with 76,000 dependents. For all individuals in the data, I observe their plan choice and administrative health insurance claims for each year during which they enroll in a plan. As is common in this type of data, I do not observe employees who choose not to enroll in a plan. Fortunately, due to the high subsidies offered by employers and the market failures present in the individual market, less than 20% of individuals forgo coverage offered to them by their employer (Kaiser Family Foundation 2013).

In order to simplify estimation, I limit the sample in the following ways. First, I only include employees who enroll no dependents. This permits me to avoid making assumptions about the family structure of the employee premium contribution which is not available in the data. Second, I limit the sample to employees who are enrolled for the full 365 days of 2006 and 2007. As described below, in order to estimate each employee's distribution of expected out-of-pocket costs, I require information on utilization and diagnoses from the year prior to plan choice. If this information is incomplete, the estimates of future costs will be biased.[16]

[15] This is the same employer used by Geruso (2013).

[16] As discussed below, in order to estimate switching costs, I require the presence of new enrollees who were not previously enrolled in a plan. Therefore, I do include enrollees not enrolled at all during the prior year. For these new enrollees, the distribution of expected out-of-pocket costs is estimated using information on utilization and diagnoses from the current year (see appendix for more details), so I require that they remain enrolled

The left columns of Table 1 shows observed characteristics of the employees in the sample. The Marketscan database includes minimal demographic information about the employees in order to protect the privacy of Truven's clients. The average age among the employees in my sample is around 41, and 60% of the employees are male. About 8% of the sample are defined as "new" employees, meaning they were not enrolled in a plan during 2006. As discussed below, this will be important for estimating switching costs. The average total annual health care costs among the employees in the sample is around $3,900.

The two PPO plans offered by this employer differ only in cost sharing, and the contracts remained constant throughout the sample period. The cost-sharing parameters of each plan are found in the left columns of Table 2. For medical costs, Plan E has a lower deductible, coinsurance rate, and out-of-pocket maximum. With respect to other costs from ER visits and prescription drugs, cost sharing is identical in the two plans.[17] The only differences between the plans are the cost sharing parameters for medical claims and the plan premiums. Unfortunately, neither the premiums nor the employee contribution to the premiums are available in the data. Because the employee contribution is a critical piece of the empirical model described below, I follow Kowalski (2013) and Geruso (2013) and estimate the contribution from the data. I discuss the estimation process in the Section 4.2.1.

3.1 Cost Model Sample

In order to estimate the choice model discussed below, I need to construct an estimate of each individual's distribution of expected costs. As discussed in Section 4.2.3 below, estimation of this distribution is likely to be more accurate with a larger dataset. Because the firm sample is relatively small, in order to estimate the cost model I augment the sample using data on 845,000 additional individuals from the Marketscan Database to form the cost model sample. The cost model sample consists of a random sample from of all individuals in Marketscan from 2006-07 enrolled in a PPO plan for at least 300 days of year t and year $t+1$. The characteristics of the cost model sample are found in column 4 of Table 1. While the means of the variables in the table differ for these two samples, as discussed below, the ranges of risk scores and ages are more relevant for comparing the samples. Additionally, the cost model sample will be validated by comparing costs predicted by the cost model to actual costs among the employees in the sample.

---

for all 365 days of the current year. This implies that the condition for remaining in the sample is that the enrollee be enrolled for either 365 days of the prior year and one month of the current year (to allow me to observe plan choice) or 365 days of the current year.

[17] All of these parameters apply only to providers in the plan's network. Claims from out-of-network providers are covered much less generously. However, there are very few of these claims, and they are not consistently and clearly identified, so I ignore them here.

4 Empirical Model

As discussed above, in order to simulate equilibrium prices and sorting under risk adjustment, I require the joint distribution of demand, total costs, and predicted costs. In this section, I will discuss how I recover each component of this distribution.

4.1 Total Costs and Predicted Costs

Because my data include the universe of health insurance claims for each individual in the sample, I observe total costs. I also observe predicted costs in the data. As discussed above, predicted costs are a linear function of individuals' risk scores. Risk scores are assigned using the following formula:

$$r_i = \frac{X_i\beta}{\bar{R}_i}$$

X_i represents a vector of "risk adjusters," or variables that predict an individual's health status, and β represents a vector of risk adjustment weights. Different risk adjustment models use different groups of variables. The models I use in the simulations include dummy variables for age-by-sex cells and a set of "Hierarchical Condition Categories," or HCCs. HCCs indicate whether an individual has a particular health condition. They are based on diagnoses found in health insurance claims. When the diagnoses are from the prior (current) year, the model is referred to as a "prospective" ("concurrent") model. HCCs are used in the models developed by HHS and CMS for the Federal Exchange, Medicare Advantage, and Medicare Part D. Details about the models studied in this paper can be found in the appendix.

For each model, the risk adjustment weights, β, were estimated using a large sample of insurance claims from the Marketscan database.[18] First, total annual costs, c_i, are regressed on the risk adjusters:

$$c_i = X_i\delta + \varepsilon_i$$

The coefficients from the regression, $\hat{\delta}$, are then normalized by dividing by the average cost in the estimation sample: $\beta = \frac{\hat{\delta}}{\bar{c}}$. This implies that an individual whose predicted cost is the average cost in the estimation sample will be assigned a risk score of 1.0.

In practice, individuals are assigned HCCs using diagnoses from their health insurance claims. Because I observe the health insurance claims for each employee in my sample, I also observe their HCCs. I calculate each individual's risk score by combining these HCCs with the pre-estimated weights for the risk adjustment models. The critical

[18] Note that I did not estimate these weights. The demographic and industry prospective and concurrent models were estimated by Verisk Health. The HHS-HCC model was estimated by HHS.

assumption here is that individuals would receive the same diagnoses in a setting with or without risk adjustment.[19]

4.2 Demand

While employees' costs can be observed in the data, their valuation of the plans is unobservable and must be estimated. Conceivably, valuation could be non-parametrically estimated by observing how employees respond to an exogenous shift in plan prices (Einav et al. 2010). However, because there is no variation in prices in my data, in order to estimate valuation, I must specify a structural model of health plan choice and use the model to estimate demand using a method similar to that used in Cohen and Einav (2007). Fortunately, the two plans at the firm I study are vertically differentiated in that they differ only in cost sharing and premiums, making the assumption that the structural model fully characterizes the employees' valuation more easily justified. I start by assuming that employees value plans based on the following von-Neuman Morgenstern expected utility function:

$$U_{ij} = \int f_{ij}(OOP)u_i(W_i, OOP, P_j, \mathbf{1}_{ij,t-1})dOOP$$

Four variables enter into the employee's utility function: $f_{ij}(OOP)$, employee i's distribution of expected out-of-pocket costs if enrolled in plan j; W_i, employee i's wealth; P_j, plan j's premium; and $\mathbf{1}_{ij,t-1}$, an indicator for whether employee i was enrolled in plan j during the previous period. The additional factor affecting individual choice is the shape of u_i. I assume that employees' preferences follow the constant absolute risk aversion (CARA) formulation. Let x_i represent the ex-post consumption level of individual i. The CARA assumption implies that

$$u_i = -\frac{1}{\gamma_i(Y_i, Z_i)} e^{-\gamma(Y_i, Z_i)x_i}$$

The shape of each individual's CARA utility function is defined by her coefficient of absolute risk aversion, γ_i, with larger values of γ_i implying higher levels of risk aversion. I define x as follows:

$$x_i = W_i - P_j - OOP + \eta(Y_i)\mathbf{1}_{jt=j,t-1} + \varepsilon_{ij}$$

The employee's consumption is a function of initial wealth, the plan premium, expected out-of-pocket costs, a switching cost incurred if the employee chooses a different plan in year t than in year $t-1$, and an i.i.d. preference shock,

[19] This assumption would be violated if plans respond to risk adjustment by "upcoding" or by increasing utilization in order to increase the number of diagnoses and thus extract a larger risk adjustment transfer (see Geruso and Layton (2014) for an example of this in Medicare). While this type of behavior is likely, it is unlikely that it would dramatically alter individuals' risk scores and the joint distribution estimated here. Additionally, if plans are identical and all individuals are equally "upcode-able," upcoding will result in the market average risk score increasing in tandem with the plan average risks scores. Because risk adjustment transfers are based on normalized risk scores, this would result in precisely the same normalized risk scores as the setting with no upcoding.

ε_{ijt}, with mean μ_ε and variance σ_ε^2.[20] Because the two plans employees of the firm can choose between vary only in their financial characteristics, I argue that this specification comes quite close to fully characterizing employee choices. With all of the components of this model, I can determine each individual's choice of plan under different levels of the price differential, P, from the model above. This will allow me to determine each individual's demand for Plan E, the third component of the joint distribution required for simulation of competitive equilibria with risk adjustment. As most of the components of the model are unobserved, they require some form of estimation. I now discuss how I estimate each component.

4.2.1 Plan Premiums: P_j

As discussed above, the data do not include any information about the employee contribution to the premium, so I must estimate it. Most employers follow a simple pricing rule based on the average cost of individuals enrolled in a plan during the prior year (see Handel (2013) and Geruso (2013)). I assume that for the firm I study, the premiums for employees without dependents are equal to the average cost among the employees enrolled in each plan during the prior year plus some loading factor, $P_t^j = AC_{t-1}^j + \alpha$. I calculate AC_{t-1}^j using the claims data from the prior year. I then assume that the employer sets the employee contributions equal to 20% of the full premium of each plan (Kaiser Family Foundation 2013). Note that for estimation of the choice model it is not important for the premiums of each plan to be accurately estimated. Instead, it is just important that the premium differential P be correct. Given the assumptions, this premium differential is:

$$P_t = 0.2(AC_{t-1}^E - AC_{t-1}^B)$$

To address the possibility that P_t is incorrectly estimated, I also include a plan specific intercept for Plan E in x_{ij}. Because all individuals pay the same prices, This intercept will capture both any idiosyncratic preference for Plan E and any bias in the estimate of P.

4.2.2 Switching Costs: η_i

There is extensive empirical evidence that individuals face substantial switching costs when choosing to move between health plans (Sinaiko and Hirth 2011, Handel 2013, Polyakova 2013). There are many reasons for this phenomenon such as the time and hassle costs of researching a new plan and switching, attachment to a network of providers, or just pure inattention or laziness.[21] Here, the source of the switching cost is unimportant. It is included in the model to

[20] The assumption of CARA utility makes W_i irrelevant, implying no income effects.
[21] See Handel (2013) for a thorough discussion of potential sources of inertia.

allow simulation of equilibrium sorting where all individuals face an active choice, as in the first year of operation of the Exchanges. In order to separate switching costs from persistent heterogeneity in preferences, I follow Handel and Kolstad (2013) by exploiting the fact that some employees in the data were not previously enrolled in a plan. While I do not observe why these enrollees are enrolling for the first time, I know that they should not face a switching cost when making their choice. To account for observable differences between new and old enrollees, I allow η to vary with observable demographic characteristics. Specifically, I assume that

$$\eta_i = \eta_0 + \eta_1 age_i + \eta_2 female_i$$

Effectively, I compare the choices of new and old enrollees with similar demographics and cost risk to estimate the switching costs. The important assumptions here are that new and old enrollees are similar with respect to unobserved variables that affect risk preferences and that there is sufficient variation in the observed characteristics (age and gender) among the new enrollees such that there is a new enrollee similar to every old enrollee. Columns 2 and 3 of Table 1 show that while the mean age is lower among the new enrollees, the range of the ages of new enrollees is almost identical to that of old enrollees.

4.2.3 Out-of-Pocket Cost Distributions: $F_{ij}(OOP)$

As discussed, the model requires an estimate of each employee's distribution of expected out-of-pocket costs in each plan, $F_{ij}(OOP)$. I construct this distribution directly from data in the cost model sample described above. To obtain a distribution of expected total costs for each employee, I first divide the full Marketscan sample into cells of employees with similar health status in year $t-1$. The cells are based on predictive measures of each individual's medical cost risk generated by sophisticated predictive modeling software developed by Verisk Health and used by health insurers and large employers to predict the costs of their enrollees. [22] The software uses information such as diagnoses and utilization found in insurance claims data from year $t-1$ to generate individual-level medical risk scores, λ_i, describing each individual's medical cost risk in year t. These risk scores are different from r_i, the risk scores used for risk adjustment described above. The risk adjustment risk scores are based only on information about diagnoses and demographics. The medical cost risk scores used here are based on the entire set of information available in the health insurance claims. This includes past utilization and spending in addition to the diagnoses and demographics used to generate the risk adjustment risk scores. In order to ensure that the estimates of each employee's cost distributions are as precise as possible, I split the sample into 500 cells based on λ_i . To ease computation, I take a random sample of

[22] I note that I sort individuals into cells based only on predictions of medical cost risk and not predictions of prescription drug cost risk because the plans offered by this employer do not vary with respect to cost-sharing for prescription drugs, making those costs irrelevant to the employee's choice. I do, however, incorporate predictions of prescription drug costs in the counterfactual simulations below.

1,000 individuals from each medical cost risk cell. For all of the individuals in a given cell, I fit a lognormal distribution with a point mass at zero to the actual medical costs of the individuals in the cell in year t and allow the lognormal parameters to vary with age and gender. The lognormal parameters plus the point mass at zero fully describe the estimates of each employee's distribution of expected total medical costs. I then use the simple cost-sharing rules for each plan to map each employee's expected total medical costs to expected out-of-pocket costs to form $F_{ij}(OOP)$.[23]

I use the cost model sample rather than the smaller choice model sample to construct $F_{ij}(OOP)$ because there is a tradeoff between cell size and the number of cells. With a larger number of cells I capture more of the private information about individuals' future costs. However, larger cells necessarily imply fewer individuals in each cell, resulting in less accurate estimation of the parameters describing $F_{ij}(OOP)$. Using the larger cost model sample avoids this tradeoff by increasing the total number of individuals in the sample. The cost of using the cost model sample rather than the choice model sample to estimate $F_{ij}(OOP)$ is the requirement of an additional assumption: Individuals in the cost model sample and the choice model sample are similar with respect to any relevant variables not used to form the cells. Given the large amount of sophisticated information used to form the cells and the large number of cells, I argue that this assumption is reasonable. Additionally, in Table 1 I show that the range of risk scores is similar for the cost model and estimation samples, implying that there are similar individuals in the samples. Table 1 also shows that for the estimation and simulation samples the average expected cost produced by the cost model is quite similar to the average realized cost among employees in the sample, implying that the estimates are not systematically biased.

4.2.4 Risk Preferences: γ_i

Each individual's coefficient of absolute risk aversion,γ_i, is the final component of the choice model. I estimate this parameter as follows. Because the two plans available to the employees in my sample differ only in cost sharing and premiums, if the employees are all risk neutral and $F_{ij}(OOP)$ and P_j are known, their optimal choices can easily be recovered by calculating the mean of each employee's distribution of expected out-of-pocket costs in each plan, adding that mean to each plan's premium, and then comparing the two sums. Whichever plan has the lower total cost (premium plus out-of-pocket costs) would be the optimal choice. Call this choice the risk-neutral optimal choice. The intuition behind the identification of γ_i is that under the assumption that $F_{ij}(OOP)$ is observed, an employee's deviation from the risk-neutral optimal choice describes her level of risk aversion (Cohen and Einav 2007). For example, if an employee faces higher total cost in Plan B than in Plan E, but she chooses Plan E anyway, she must be risk averse, and the size of the cost difference identifies the extent of her risk aversion. This method is also used by Handel (2013) and

[23] As mentioned above, there is no cost-sharing for ER and preventive visits. Because the prediction software predicts total medical risk, rather than medical risk other than ER and preventive visits, I cannot separately predict an employee's use of these services. Instead, I ignore the fact that they are priced differently and combine them with other medical spending. In practice, these services make up a relatively small portion of spending and have a small impact on inferences about the distribution of $F_{ij}(OOP)$

Geruso (2013).

To ensure the joint distribution of demand, total costs, and predicted costs is fully characterized, I allow γ_i to vary with a set of demographic variables, Y_i, along with employees' total ex-post realized costs, c_i, and risk adjustment risk scores (predicted costs), r_i, to allow for heterogeneity in risk preferences. Specifically, I assume that γ_i can be described as follows:

$$\gamma_i(Y_i, Z_i) = \beta_0 + \beta_1 age_i + \beta_2 ln(\lambda_i) + \beta_3 age * ln(\lambda_i) + \beta_4 ln(r_i)$$

Allowing γ_i to vary with c_i is motivated by previous research that has shown that the correlation between risk preferences and cost risk can influence the degree and direction of selection in equilibrium (see Finkelstein and McGarry (2006), Cohen and Einav (2007), Einav et al. (2013), and Handel et al. (2013)). The inclusion of r_i in the risk preference equation is motivated by the graphical analysis above. Recall that the equilibrium consequences of risk adjustment depend on the relationship between demand and predicted costs. Thus, in order to accurately simulate equilibrium prices and sorting under risk adjustment, it is critical that the model fully capture this relationship. Because demand is a function of γ_i, it is necessary to allow γ_i to vary by the risk scores that determine predicted costs. In fact, if this correlation is not allowed for in the estimation of γ_i, the counterintuitive situation where risk adjustment results in a larger price differential and fewer individuals enrolling in Plan E would be missed in the simulations.

4.2.5 Limitations

While this demand specification characterizes the choices of consumers quite nicely, it does rely on a few important assumptions. First, I assume that when making their choices between the two plans, employees are using the same distribution of expected out-of-pocket costs that I assign to them. While it is possible that individuals know more than what I am able to predict, it is unlikely that they know much more. On the other hand, it is also possible that individuals know much less than the model suggests. To deal with this problem, I perform two sets of robustness checks using alternative specifications of $F_{ij}(OOP)$. In the first set of robustness checks, I test the sensitivity of my results to forming $F_{ij}(OOP)$ using 50 or 1000 cells in λ_i instead of 500. In the second set, I allow consumers to have private information by forming a new distribution of expected costs, $G_{ij}(OOP) = \alpha oop_{ij} + (1-\alpha)F_{ij}(OOP)$, where oop_i is i's out-of-pocket cost in plan j based on i's actual realized costs. I set α equal to .05, .10, .15, .20, and .25. This still leaves open the possibility that rather than just using limited information in a rational manner, individuals actually use sophisticated information but they do so irrationally. There is evidence for this type of behavior (Handel and Kolstad 2013, Abaluck and Gruber 2011). It is important to note, however, that "mistakes" will likely be captured as risk preferences in this model. This is not a problem for using the joint distribution of demand, total costs, and predicted costs to simulate competitive equilibria with and without risk adjustment. However, it does present a problem for inference about the welfare consequences of risk adjustment because the area below the demand curve and above the incremental marginal cost curve may not actually represent consumer surplus (Handel et al. 2014).

4.2.6 Estimation

I estimate the parameters of the model using a simulated maximum likelihood approach similar to the method used by Handel (2013) and Geruso (2013) and outlined in Train (2009). Estimation begins by fixing the parameter vector, Φ, and taking a draw from the distribution the error term, ε^s, which is assumed to follow a normal distribution. Next, Q draws are taken from each employee's expected total cost distribution and run through each plan's cost-sharing parameters to simulate $F_{ij}(OOP)$ for each employee i and plan j. Each employee's expected utility from plan j is then estimated by averaging the CARA utility function over the Q draws from $F_{ij}(OOP)$:

$$E[U_{ij}] = \sum_{q=1}^{Q} e^{-\gamma_i(Y_i,Z_i)(P_j+OOP_{ij}^q+\eta(Y_i)\mathbf{1}_{ij,t-1}+\varepsilon_{ij}^s)}$$

where

$$\begin{aligned} \gamma_i(Y_i,Z_i) &= \beta_0+\beta_1 age_i+\beta_2 ln(\lambda_i)+\beta_3 age*ln(\lambda_i)+\beta_4 ln(r_i) \\ \eta(Y_i) &= \eta_0+\eta_1 age_i+\eta_2 female_i \end{aligned}$$

Given the draw of the error term ε^s and the expected utility estimates for each plan they imply, I could simulate each employee's choice by comparing the expected utilities for Plan E and Plan B and assigning employee i to the plan with higher expected utility, i.e. an accept-reject simulator. However, the accept-reject simulator can cause problems in the estimation process due to flat portions of the likelihood function where no employees choose to move from one plan to the other and undefined portions of the log-likelihood function due to some individuals having a zero probability of enrolling in one of the plans. Both of these problems could potentially be solved by including a larger number of draws of the random components; however, for both issues to be fully resolved, the number of necessary draws would approach infinity. Instead, I use a smoothed accept-reject simulator . The simulator I use was developed by Handel (2013) and simulates the probability that the employee will choose each plan according to the following function:

$$R_{ij}^s(j=j^*) = \frac{\left(\frac{\frac{1}{-E[U_{ij^*}^s]}}{\sum_J \frac{1}{-E[U_{ij}^s]}}\right)^\tau}{\sum_{\hat{j}}\left(\frac{\frac{1}{-E[U_{ij}^s]}}{\sum_J \frac{1}{-E[U_{ij}^s]}}\right)^\tau}$$

The form of the simulator ensures that the estimated probability increases (decreases) when the accept-reject simulator would increase (decrease) and that the probability lies between zero and one. As the smoothing parameter, τ, becomes large, the simulator becomes identical to the accept-reject simulator.

The probability of choosing each plan is calculated for each draw, s, of the the error term. To find the simulated probability that the employee chooses plan j, given the draw from the parameter distribution, I average over the smoothed values,

$$P_{ij}(\Phi) = \frac{1}{S}\sum_{s=1}^{S} R_{ij}^{s}$$

The simulated log-likelihood function is then defined as

$$SLL(\Phi) = \frac{1}{N}\sum_{i=1}^{N}\sum_{j=1}^{J} d_{ij} ln P_{ij}(\Phi).$$

d_{ij} is equal to one if the employee actually chose plan j and zero otherwise. The likelihood function is quite intuitive in that it will achieve its maximum where the simulated probability that each employee chose plan j is as close to one as possible when the employee actually chose plan j and as close to zero as possible otherwise. I use standard numerical techniques to find the parameter vector, Φ, that maximizes the likelihood function. In the actual estimation I set $Q = 50$, $S = 75$, and $\tau = 2$.

4.2.7 Model Results

Table 3 presents the results from the choice model discussed above. The estimates of the coefficient of absolute risk aversion are similar to estimates in the health insurance literature (Geruso 2013, Handel 2013). The average estimate of the coefficient of absolute risk aversion is $6.3 * 10^{-4}$. To aid interpretation, this level of risk aversion implies that the average employee in the sample would be indifferent between the status quo and a gamble where she will win \$100 or lose \$94 with equal probability. Perhaps more importantly, the estimates imply that risk aversion is negatively correlated with total costs, negatively correlated with prospective risk adjustment risk scores, and positively correlated with concurrent risk adjustment risk scores. If taken literally, this implies that controlling for prospective and concurrent risk scores, employees with high levels of total cost have lower levels of risk aversion. This result is similar to the finding of Handel et al. (2013) and is consistent with the notion of the “worried well.”

The average estimated switching cost is close to \$2,500. While this is quite large, it is consistent with other estimates in the literature (Handel 2013, Handel and Kolstad 2013). To put this in context, this is over 60% of the total average health care costs among the employees in the sample. Such large switching costs imply that simulations of active and passive choices will produce substantially different results.

Given the assumption of CARA utility and the estimates of employee risk aversion (γ_i) and the expected out-of-pocket cost distribution ($f_{ij}(OOP)$), I can simulate individuals' choices between Platinum and Bronze plans in an Exchange setting. This provides the final component of the joint distribution of total costs, predicted costs, and demand.

5 Counterfactual Simulations

In order to simulate competitive equilibria with and without risk adjustment, I first expand the sample to form a new simulation sample. The simulation sample includes all single-coverage employees from 2005-07. The sample is restricted in the same ways as the choice model sample described above (i.e. must be enrolled for all 365 days of the year, etc.). For this sample, total costs and risk scores for each individual are calculated or estimated as described in Section 4. Summary statistics for this simulation sample can be found in the final column of Table 1.

I simulate competitive equilibria with and without risk adjustment in an environment similar to the Exchanges currently being established throughout the United States. Specifically, each individual is required to choose either a Bronze plan or a Platinum plan where the two plans are vertically differentiated in that they only differ in cost sharing and premiums. The Platinum plan is much more comprehensive than the Bronze plan. Plan cost sharing is described by standard non-linear price schedules where individuals pay the full cost of all care received up to a deductible, then some portion of each additional dollar of care up to an out-of-pocket maximum. The exact price schedules can be found in columns 3 and 4 of Table 2. The simulated Platinum (Bronze) plan has a deductible of $0 ($4500) a coinsurance rate of 20% (20%) and an out-of-pocket maximum of $1500 ($6500). The cost-sharing parameters were chosen using the 2014 version of the actuarial value calculator provided by the Department of Health and Human Services.[24]The choice to use the most and least comprehensive plans available on the Exchange in the simulation was deliberate. In order to capture the most accurate picture of adverse selection on the Exchange, the simulation should include these two options.[25]

One additional component is required to complete the joint distribution of demand, total costs, and predicted costs. Demand is a function of two factors: risk preferences and the distribution of expected out-of-pocket costs. Risk preferences are assigned to each individual according to the estimated parameters from the model above, found in Table 3. The distributions of expected out-of-pocket costs for the new Platinum and Bronze plans are estimated using techniques similar to those described in Section 4. The key difference is that under the Platinum and Bronze

[24]This is the calculator that all insurers are required to use to ensure that their plans meet the actuarial value requirements of the ACA. The calculator can be found at http://www.cms.gov/CCIIO/Resources/Regulations-and-Guidance/Downloads/av-calculator-final.xlsm

[25]In practice, most consumers will choose between 4 tiers of plans with varying levels of comprehensiveness. However, the competitive equilibrium in this environment would be quite complex and extremely difficult to model while not providing much additional intuition.

plans, cost sharing for prescription drug utilization is not assumed to be identical in the two plans. Instead, the cost-sharing parameters described above are applied to total health care costs, the combination of prescription drug costs and medical costs. Therefore, in order to construct the distributions of expected out-of-pocket costs for the Platinum and Bronze plans, I first need to estimate each individual's distribution of expected *total* health care costs, rather than the distributions of expected total medical costs used to estimate the choice model. I do this by dividing the cost model sample discussed above into 500 cells based on total cost risk instead of medical cost risk, where total cost risk is generated using a sophisticated predictive model analogous to the one used to generate medical cost risk. I then follow the procedure outlined above to complete the construction of the expected out-of-pocket cost distributions, here using the Platinum and Bronze cost-sharing parameters rather than the Enhanced and Basic plan parameters used above. The combination of risk preferences and the expected out-of-pocket cost distribution allow me to calculate each individual's expected utility under plan j by taking Q draws from the estimated $F_{ij}(OOP)$ and using the following expression

$$E[U_{ij}] = \sum_{q=1}^{Q} e^{-\gamma_i(P_j + OOP_{ij}^q)}$$

I can then determine which plan i will choose given a price differential P, providing the final component of the joint distribution: Demand.[26]

5.1 Correlations

As shown in Section 2, the effect of risk adjustment on prices is revealed by the correlations between demand for the Platinum plan and total costs and demand for the Platinum plan and predicted costs. If both correlations are positive, then risk adjustment will result in a smaller price differential and higher enrollment in the Platinum plan. Before, moving to the simulations of equilibrium with and without risk adjustment, I present these correlations. Figure 6 shows the correlation between demand for the Platinum plan and total costs. Quantiles of willingness-to-pay for Platinum over Bronze are on the x-axis, and a normalized measure of total costs (total cost divided by average cost) is on the y-axis. It is clear that the correlation is positive, i.e. total cost is increasing in willingness-to-pay. This implies that as the price differential decreases, the marginal Platinum enrollee will have lower total cost than the average Platinum enrollee, i.e. the Platinum plan will be adversely selected.

Figure 7 shows the correlations between demand and the risk scores from each of the types of risk adjustment simulated below. I show the correlation between demand and the risk score rather than demand and predicted costs

[26] Note that in the simulations consumers are assumed to pay the full incremental cost of enrolling in the Platinum plan rather than the subsidized cost assumed in the choice model. The subsidized cost was used to estimate the choice model because it is likely to correspond closely with the price the employees of the firm actually faced. The full incremental cost is used in the simulations because this is the price that will be faced by individuals purchasing coverage through the Exchanges.

because predicted costs are a linear function of the risk score, implying that the correlations will be identical. It is clear that the correlation between demand for the Platinum plan and each type of risk score is also positive. This implies that as the price differential decreases, the marginal Platinum enrollee will have a lower risk score than the average Platinum enrollee, i.e. the Platinum plan will be adversely selected on the risk score. The model and figures in Section 2 show that because the correlations between demand and total costs and demand and predicted costs are both positive, in this setting risk adjustment will result in a smaller price differential and increased enrollment in the Platinum plan.

It is interesting to note that the correlation between demand and predicted costs is quite different for demographic-based risk adjustment (Figure 7A) than for the 3 types of diagnostic-based risk adjustment (Figures 7B-C). Comapred to diagnostic-based risk adjustment, demographic-based risk adjustment does a relatively poor job of capturing the costs of the highest cost/highest demand enrollees. It is also interesting to note that the correlations between predicted costs and demand are almost identical for each of the three types of diagnostic-based risk adjustment. While the two concurrent models (Figures 7C and 7D) explain around 40% of the variance in total costs, the prospective model explains less than 20%. Figure 7B reveals, however, that despite explaining a much smaller portion of the variance in total costs, the prospective model seems to capture almost all of the variation in total costs that is correlated with demand.

5.2 Welfare

In order to assess the welfare consequences of risk adjustment in the simulations, I follow Einav et al. (2010) and use a certainty equivalent concept. The certainty equivalent, e_{ij}, is defined as the value, e, that makes individual i indifferent between paying e and facing the uncertain loss under insurance plan j. I calculate the certainty equivalent for individual i under plan j by finding the value of e_{ij} that makes the following expression true

$$-\frac{1}{\gamma_i}e^{-\gamma_i e_{ij}} = \int e^{\gamma(P_j+OOP_{ij})}dF_{ij}(OOP).$$

The certainty equivalent is convenient because it provides a way to determine i's valuation of the insurance plan in dollars. Given an individual's certainty equivalent under each plan, I can calculate individual i's willingness-to-pay for the Platinum plan in dollars by subtracting e_{iB} from e_{iP}: $WTP_{iP} = e_{iP} - e_{iB}$. The willingness-to-pay for the Platinum plan incorporates the difference in out-of-pocket costs paid by i and the difference in uncertainty under the two plans. It also represents consumer surplus from i moving from the Bronze plan to the Platinum plan.

Total welfare, however, must also account for changes in the costs paid by the insurer. Because the difference in out-of-pocket costs under the two plans just represents a transfer from the insurer to the consumer, it does not affect total welfare. Only the decreased uncertainty will impact total welfare. Thus, the change in total welfare from i moving from the Bronze plan to the Platinum plan is

$$\Delta W_i = WTP_{iP} - (c_i^P - c_i^E)$$

where c_i^j is the plan cost from enrolling i in plan j, making $(c_i^P - c_i^E)$ i's incremental marginal cost. ΔW_i represents the incremental welfare improvement from moving i from the Bronze to the Platinum plan. This implies that the change in total welfare resulting from a move from a setting with no risk adjustment to a setting with risk adjustment is equal to

$$\Delta Welfare = \sum_{i=1}^{N} \Delta W_i \mathbf{1}(E[U_{iP}^{RA}] > E[U_{iB}^{RA}]) - \sum_{i=1}^{N} \Delta W_i \mathbf{1}(E[U_{iP}^{NRA}] > E[U_{iB}^{NRA}])$$

where $E[U_{ij}^{RA}]$ is i's expected utility under plan j given the equilibrium prices with risk adjustment and $E[U_{ij}^{NRA}]$ is i's expected utility under plan j given the equilibrium prices without risk adjustment. The intuition for this measure is that welfare only changes when individuals move from one plan to another, and when an individual moves, welfare either increases or decreases by ΔW_i, depending on whether i is moving from Bronze to Platinum or Platinum to Bronze.

5.3 Equilibrium

In order to simulate equilibrium prices and sorting in this environment, I first need to establish an equilibrium concept. Handel et al. (2013) show that in this setting, the competitive equilibrium can be found using the following algorithm where P^P and P^B represent the premiums of the Platinum and Bronze plans, $P = P^P - P^B$, $AC^P(P)$ and $AC^B(P)$ represent the average plan costs of enrollees in the Platinum and Bronze plans given price differential P, AC^E and AC^B represent the average plan costs of the entire population in the Platinum and Bronze plans, and $\Delta AC(P) = AC^P(P) - AC^B(P)$:

1. If $\Delta AC(P) < P \; \forall P$ then the entire market enrolls in the Platinum plan and $P^P = AC^E$
2. If $\exists P$ such that $P = \Delta AC(P)$ then the equilibrium value of P is equal to $P^* = \min(P : P = \Delta AC(P))$ and consumers sort according to willingness-to-pay
3. If $\Delta AC(P) > P \; \forall P$ then the entire market enrolls in the Bronze plan and $P^B = AC^B$

Figure 8 illustrates the equilibrium search with no risk adjustment. P is on the x-axis. The light blue line represents the 45-degree line, and the orange line represents $\Delta AC(P)$. If there is an interior equilibrium, it will be where $\Delta AC(P) = P$.

It is clear that there is no interior equilibrium in this setting and that $\Delta AC(P) > P$ for all values of P. This implies that in equilibrium, the entire market enrolls in the Bronze plan and $P^B = AC^B$. This is also known as market unraveling and is the same result found by Handel et al. (2013). In this setting, the correlation between demand and total costs is so strong that there is no price differential at which any part of the market enrolling in the Platinum plan would result in a competitive equilibrium where both plans earn zero profits.

5.3.1 Risk Adjustment

With risk adjustment, the equilibrium concept remains the same, but the relevant plan average costs change. Risk adjustment is implemented by assuming the regulator gives each plan the following transfer

$$T_j(P) = \left(\frac{\bar{R}_j}{\bar{R}} - 1\right)\bar{P}$$

where $\bar{R}_j$ represents the average risk score of the enrollees in plan j, $\bar{R}$ represents the average risk score in the entire market, and $\bar{P}$ represents the average premium in the market. This is analogous to the individual-level risk adjustment framework developed in Section 2. With risk adjustment, equilibrium is where the premium differential is equal to the incremental average risk adjusted cost, $\Delta AC_{RA}(P)$. The algorithm for finding the competitive equilibrium remains the same, except $\Delta AC(P)$ is replaced with $\Delta AC_{RA}(P)$. Note that $AC^j = AC^j_{RA}$ for both plans because when the entire market is enrolled in the same plan $T_j = 0$.

I simulate four types of risk adjustment: demographic based on age-by-sex cells, industry prospective and concurrent diagnostic based on HCCs, and the HHS-HCC model being implemented in the Exchanges.[27] Figure 9 illustrates the equilibrium search under these four forms of risk adjustment. Again, the light blue line represents the 45-degree line and the orange line represents $\Delta AC(P)$. The dark blue, gray, gold, and green lines represent $\Delta AC_{RA}(P)$ under demographic, prospective, concurrent, and HHS-HCC risk adjustment, respectively. All four risk adjustment policies shift the incremental average cost curve down. Interestingly, none of the risk adjustment policies has much of an effect on the right tail of the curve. This suggests that these policies do a relatively poor job of explaining the costs of the individuals who value the Platinum plan the most: the sick. Under demographic risk adjustment, $\Delta AC_{RA}(P) > P$ for all values of P and the entire market still enrolls in the Bronze plan. In other words, demographic risk adjustment has no effect on equilibrium prices or sorting in this setting. However, under prospective, concurrent, and HHS-HCC

[27] The transfer for HHS-HCC risk adjustment is slightly different from the formula used for the other forms of risk adjustment. It is

$$T_j(P) = \left(\frac{\bar{R}_j}{\bar{R}} - \frac{AV_j}{\bar{AV}}\right)\bar{P}$$

where AV_j is the actuarial value of plan j and $\bar{AV}$ is the enrollment-weighted average actuarial value in the market. The actuarial value adjustment is due to the fact that the HHS-HCC risk scores, unlike the other risk scores, increase with the actuarial value of the plan.

risk adjustment $\Delta AC_{RA}(P)$ crosses the 45-degree line, implying that there exists an interior equilibrium. With all three forms of diagnostic risk adjustment there are in fact multiple points where $P = \Delta AC_{RA}(P)$. Recall that according to the algorithm, the competitive equilibrium value of P, P^*, is the smaller value of P for which $P = \Delta AC_{RA}(P)$. According to the algorithm then, under prospective risk adjustment $P^* = \$1,031$, under concurrent risk adjustment $P^* = \$1,050$, and under HHS-HCC risk adjustment $P^* = \$1,806$.

Figure 10 illustrates the equilibrium allocations of individuals across plans in a more familiar way, similar to the figures in Section 2. In this figure, enrollment in the Platinum plan is on the x-axis. Again, the orange, dark blue, gray, gold, and green lines represent the incremental average cost curve under no risk adjustment, demographic risk adjustment, prospective risk adjustment, concurrent risk adjustment, and HHS-HCC risk adjustment, respectively. The light blue line reflects demand or willingness-to-pay for the Platinum plan relative to the Bronze plan. For prospective, concurrent, and HHS-HCC risk adjustment, the equilibrium price P^* is at the smaller of the two points where the gray, gold, and green incremental average risk adjusted cost curves cross the light blue demand curve. The equilibrium price and Platinum enrollment under HHS-HCC risk adjustment are highlighted with the black lines lines. Recall that with no risk adjustment or demographic risk adjustment, the entire market enrolls in the Bronze plan. The figure shows that under prospective, concurrent, and HHS-HCC risk adjustment, a substantial portion of the market, over 60% for HHS-HCC risk adjustment and over 80% for the others, will enroll in the Platinum plan.

The results of these simulations can also be found in Table 4. The table clearly shows that diagnostic risk adjustment compresses the premiums of the Platinum and Bronze plans and undoes a substantial portion of market unraveling. Changes in welfare due to risk adjustment are also found in Table 4. The welfare calculations suggest that individuals in this market would place a high value on diagnostic risk adjustment, over $600 per person per year for HHS-HCC risk adjustment and more than $700 per person per year for the others. This suggests huge welfare gains from risk adjustment, around 20% of average total health care costs in this population.

It is also interesting to note that the effects of prospective and concurrent risk adjustment are quite similar, despite the concurrent model explaining a substantially larger portion of individuals' total costs. This is largely due to the fact that there is little correlation between demand and the extra dimensions of total costs explained by the concurrent model but not by the prospective model. This is an important and fairly intuitive finding. If the extra costs explained by the concurrent model are unpredictable acute costs, they are unlikely to affect an individual's plan choice and thus unlikely to be correlated with demand.

5.3.2 Reinsurance

During the first three years, the Exchanges will feature a reinsurance policy as well as HHS-HCC risk adjustment. The reinsurance policy reimburses health plans for 80% of an enrollee's plan costs above a threshold of $60,000 and below a cap of $250,000. Insurers are expected to purchase private coverage for costs exceeding $250,000, with a coinsurance rate of 85%. Reinsurance is also used in the Medicare Part D program. Recall that Figures 6 and 7 showed that diagnostic risk adjustment flattens the incremental average cost curve for all but the sickest individuals in the market. This suggests that when risk adjustment is combined with reinsurance, adverse selection problems could be reduced even further.

To explore this possibility I simulate equilibrium prices and sorting with each form of risk adjustment combined with reinsurance. I simulate reinsurance by assuming that for each enrollee, plans receive a payment equal to 85% of any plan costs exceeding $60,000 within a year. The simulated reinsurance program is funded with a uniform actuarially fair per capita premium equal to the expected per capita reinsurance payment in the market. The equilibrium results with reinsurance are found in Table 4 and in Figure 11. The figure and table show that even with reinsurance, when there is no risk adjustment or demographic risk adjustment, the market still fully unravels and everyone enrolls in the Bronze plan. With diagnostic risk adjustment, however, the incremental average risk adjusted cost curves are even flatter than they were without reinsurance. The flattening of the curves is especially apparent for the highest cost/highest willingness-to-pay enrollees at the far right of the figure.[28] The results in the table indicate that when concurrent risk adjustment is combined with reinsurance, market unraveling is entirely undone, with the premium differential shrinking enough to induce 100% of market participants to enroll in the Platinum plan. Similarly, with prospective and HHS-HCC risk adjustment, result in 98% and 82% of enrollees choosing the Platinum plan, respectively. However, the additional welfare gains from reinsurance when combined with prospective and concurrent risk adjustment, around $20-$30, while non-trivial are small relative to the gains from diagnostic risk adjustment. The additional gains from reinsurance when combined with HHS-HCC risk adjustment are more substantial, around $150. This finding complements the finding of Zhu et al. (2014) that when combined with risk adjustment, reinsurance can substantially weaken plans' incentives to cream skim healthy enrollees.

5.3.3 Age-based Pricing

The ACA allows for premiums in the Exchanges to vary by age as long as they remain within a 3:1 ratio. Because age-based pricing causes the prices paid by enrollees to be closer to their expected costs, this policy has the potential to undo some of the adverse selection problems resulting from uniform pricing that I study in this paper. This suggests that the effects of risk adjustment may not be as large in practice due to this regulation. In fact, Shi (2013) finds that welfare in a simulated Exchange is highest when risk adjustment is combined with age-based prices. To study this question, I implement age-based pricing as it is implemented in the Federal Health Insurance Exchange. In the Federal Exchange, plans submit one price that applies to a 21 year old. An individual's price is this price multiplied by an

[28] It is interesting that for the diagnostic-based risk adjustment policies, Figure 9 indicates that after a certain point $\Delta AC_{RA}(P)$ begins to fall. There are two potential reasons for this. First, the risk adjustment model does not take reinsurance into account, so the plan is effectively being reimbursed twice for the highest cost individuals. Second, at the far right side of the figure, there are only a few individuals left in the Bronze plan, making $\Delta AC_{RA}(P)$ quite sensitive to the random component of costs.

age-specific weight assigned by HHS. For example, the weight for a 25 year-old is 1.004, the weight for a 40 year-old is 1.278, and the weight for a 64 year-old is 3.0.[29]In practice, risk adjustment combined with mandated age-curve may "over-compensate" for costs correlated with age and cause plan revenues for an individual to be less correlated with costs than without the mandated age-curve.

The results from simulations including the HHS-mandated age-curve can be found in Table 5. Interestingly, the results are largely unchanged from the uniform price case. The bolded case in Table 5 with the age-curve, reinsurance, and HHS-HCC risk adjustment is the full policy being implemented in the Exchanges. It is clear that this policy goes a long way toward undoing the problems cause by adverse selection with 82% of the market enrolling in the Platinum plan, resulting in a welfare gain of over $700.

5.3.4 Robustness

As discussed in Section 4.2.5, the results of the choice model estimation and the simulations depend on the assumption that $F_{ij}(OOP)$ is specified correctly. Up until this point, I have assumed that individuals can predict to which of 500 cells in λ_i they belong. I now test the robustness of the results to using 50 or 1,000 cells instead of 500. The choice model results with 50 and 1,000 cells can be found in Table A1 in the appendix. It is clear that the parameter estimates change very little as I vary the number of cells. The mean coefficient of absolute risk aversion is increasing in the number of cells, but only slightly. The mean switching cost, on the other hand, is decreasing in the number of cells, but again by a relatively small amount.

Despite the similar parameter estimates, it is possible that changing the number of cells could have large effects on the simulation results. This is due to the fact that $F_{ij}(OOP)$ plays a large role in individuals' choices in the simulations, so changing it could influence the equilibrium results. Equilibrium prices and sorting with 50 and 1,000 cells can be found in Tables 6 and 7. Results are shown with and without reinsurance. Equilibrium prices and sorting are virtually identical with 50, 500, or 1,000 cells.

Despite the fact that λ_i incorporates a great deal of information from claims data into a sophisticated prediction of each individual's costs, it may be the case that some individuals have additional private information about their future spending. For example, a woman who intends to get pregnant during the next year may not have anything in her prior insurance claims suggesting her intention. In this case, the individual will know more about her future costs than λ_i

[29]This is different from the age-based pricing studied by Shi (2013) in that HHS forces insurers to use their age curve rather than allowing insurers to set their own age-based prices according to age-specific expected costs. I choose to use the fixed age-curve approach because this is the approach used in the Exchanges in every state. While the alternative approach is interesting, and was used in Massachusetts prior to the ACA, the concept of equilibrium is much more complex due to multiple age-based risk pools and is beyond the scope of this paper.

would suggest. To allow for this possibility, I allow individuals to have 5%, 10%, 15%, 20%, or 25% of $F_{ij}(OOP)$ based on their actual realized costs in the next year. The choice model estimates incorporating private information are found in Table A2. Again, the parameter estimates vary little across specifications.

Equilibrium prices and sorting with 10% and 20% private information can be found in Tables 8 and 9. Interestingly, when there is just a little private information, prospective and HHS-HCC risk adjustment do not undo market unraveling. Under these policies, the entire market enrolls in the Bronze plan. With concurrent risk adjustment, however, a substantial portion of market unraveling is still undone, with 85% of the market enrolling in the Platinum plan. This is quite different from the case with no private information where the effects of prospective and concurrent risk adjustment were virtually identical.

The reason for this result is a subtle but important one. Recall that in order to undo the consequences of adverse selection, the correlation between predicted costs and demand needs to mimic the correlation between total costs and demand. With no private information, the correlation between total costs and demand is based entirely on the portion of costs that are predictable using λ_i. Prospective risk adjustment do quite well at mimicking the correlation between demand and λ_i because, like λ_i, the prospective risk scores are based entirely on information (diagnoses) from health insurance claims from the prior year. Concurrent risk adjustment, however, is based on information from health insurance claims from the current year, allowing it to capture a portion of the private information that is used by individuals to choose between the plans and is thus correlated with demand. This implies that the incremental benefit from concurrent over prospective risk adjustment is likely increasing in the amount of private information used by individuals in plan choice. It is interesting to note, however, that despite being a concurrent risk adjustment model, the HHS-HCC risk adjustment policy is not able to undo market unraveling.

Tables 8 and 9 also show equilibrium when reinsurance is combined with risk adjustment and individuals have private information about their future health care costs. Interestingly, the results here are similar to the case with no private information. This implies that reinsurance captures a substantial portion of the correlation between the private information and demand. This result also suggests that the finding that when combined with reinsurance risk adjustment is quite powerful at undoing the effects of adverse selection is robust to assumptions about the structure of individuals' distributions of expected costs.

6 Discussion

The tradeoff between choice and adverse selection is a recurring theme among health economists. In the absence of choice, there is no potential for adverse selection because the costs of all consumers are combined in one risk pool. However, there is also no potential for efficiency gains from competition or from accommodating preference heterogeneity. Risk adjustment presents an opportunity to limit adverse selection problems by pooling a portion of consumers' costs across plans while still allowing choice and competition. Even the imperfect forms of risk adjustment studied in this paper appear to be able to eliminate a substantial portion of the welfare loss caused by adverse selection in a competitive environment similar to the Exchanges.

While the results in this paper are compelling, they are limited by some important caveats. First, the estimates of consumer preferences used in the simulations are based on a highly parametric structural model. Because the data do not include any premium variation (or even the premiums themselves!), I am unable to non-parametrically estimate an individual's willingness-to-pay for insurance. While the setting in which the individuals choose plans is quite simple and potentially easily characterized, the process by which consumers make choices in the real world is complex. There is evidence of important behavioral frictions in health plan choice (Abaluck and Gruber 2011, Handel 2013, Handel and Kolstad 2013), and although I attempt to control for perhaps the most important of these, inertia, in my estimation, it is controlled for imperfectly and makes up only one of many potential frictions. Additionally, my sample of employees is probably not representative of the individuals likely to enroll in plans through the Exchanges. The sample is probably higher income and more risk averse than the low income and currently uninsured individuals likely to enroll in an Exchange plan.

Additionally, throughout this study, the assumptions of a strong mandate and no moral hazard were maintained. If there were moral hazard, equilibrium pricing and sorting results would likely be similar, but the welfare consequences of risk adjustment would not be so extreme. This is due to the fact that with moral hazard it would not be optimal for the entire market to enroll in the Platinum plan. This possibility is related to the discussion of moral hazard in Einav and Finkelstein (2011). Additionally, if the mandate is weak and consumers can opt out of the market, the equilibrium consequences of risk adjustment could be quite different. As shown in the simulations, risk adjustment raises the premium of the Bronze plan. If the mandate is weak, this could easily result in healthy consumers dropping out of the market entirely, potentially resulting in the entire market unraveling. This issue is beyond the scope of this paper, but presents a promising area for future research. Here, I just point out that a large portion of the individuals participating in the Exchanges will be receiving subsidies that are based on the price of the second-cheapest Silver Plan. This implies that for a large segment of the market, the absolute prices of the Bronze and Silver Plans don't actually matter; all that matters is the difference between the Bronze or Silver Plan price and the Platinum plan price. This is what I simulated in this paper. However, because subsidies are based on the price of the second-cheapest Silver Plan, and the simulations in this paper show that risk adjustment is likely to substantially increase the equilibrium price of that plan, risk adjustment may result in larger subsidy payments than would otherwise be incurred, resulting in the program not being "budget neutral" as policymakers claim.

Despite the potential limitations, however, consumer choice is likely to follow similar patterns to those shown in the counterfactual simulations. Therefore, the effect of risk adjustment on equilibrium prices, sorting, and welfare while perhaps not perfectly estimated, is likely to be substantial. This is true not only in the Exchanges but also in Medicare Part D where premiums are set competitively, some portions of the contracts are fixed, and demand is correlated with predicted costs (Polyakova 2014). When combined with its potential beneficial effects on cream-skimming problems (Glazer and McGuire 2000), this makes risk adjustment a powerful tool for ameliorating adverse selection problems and improving welfare within competitive insurance markets. As these markets mature and data becomes available, it will be interesting to observe the correlations between demand, costs, and predicted costs and the effects of risk adjustment in practice.

References

[1] Jason Abaluck and Jonathan Gruber. Choice inconsistencies among the elderly: Evidence from plan choice in the medicare part d program. *American Economic Review*, 101:1180–1210, 2011.

[2] George A. Akerlof. The market for "lemons": Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84:488–500, 1970.

[3] Jason Brown, Mark Duggan, Ilyana Kuziemko, and William Woolston. How does risk selection respond to risk adjustment? evidence from the medicare advantage program. 2011.

[4] M. Kate Bundorf, Jonathan Levin, and Neale Mahoney. Pricing and welfare in health plan choice. *American Economic Review*, 102(7):3214–48, 2012.

[5] Alma Cohen and Liran Einav. Estimating risk preferences from deductible choice. *American Economic Review*, 97:745–788, 2007.

[6] David M. Cutler and Sarah J. Reber. Paying for health insurance: The trade-off between competition and adverse selection. *The Quarterly Journal of Economics*, 113:433–466, 1998.

[7] Liran Einav and Amy Finkelstein. Selection in Insurance Markets: Theory and Empirics in Pictures. *Journal of Economic Perspectives*, 25(1):115–38, Winter 2011.

[8] Liran Einav, Amy Finkelstein, and Mark R. Cullen. Estimating Welfare in Insurance Markets Using Variation in Prices. *The Quarterly Journal of Economics*, 125(3):877–921, 2010.

[9] Randall P. Ellis and Thomas G. McGuire. Predictability and predictiveness in health care spending. *Journal of Health Economics*, 26:25–48, 2007.

[10] Amt Finkelstein and Kathleen McGarry. Multiple dimensions of private information: Evidence from the long-term care insurance market. *American Economic Review*, 96:938–958, 2006.

[11] Kaiser Family Foundation. 2013 employer health benefits survey, 2013.

[12] Michael Geruso. Selection in employer health plans: Homogeneous prices and heterogeneous preferences. 2012.

[13] Michael Geruso and Timothy J. Layton. Risk selection, risk adjustment, and manipulable medical coding: Evidence from medicare. 2014.

[14] Jacob Glazer and Thomas G McGuire. Optimal risk adjustment in markets with adverse selection: an application to managed care. *The American Economic Review*, 90(4):1055–1071, 2000.

[15] Jacob Glazer and Thomas G McGuire. Setting health plan premiums to ensure efficient quality in health care: minimum variance optimal risk adjustment. *Journal of Public Economics*, 84(2):153–173, 2002.

[16] Jacob Glazer and Thomas G. McGuire. Gold and silver plans: Accomodating demand heterogeneity in managed competition. *Journal of Health Economics*, 30:1011–9, 2011.

[17] Jacob Glazer, Thomas G. McGuire, and Julie Shi. Risk adjustment of health plan payment to correct inefficient plan choice from adverse selection. 2013.

[18] Benjamin Handel. Adverse selection and inertia in health insurance markets: When nudging hurts. *American Economic Review, forthcoming*, 2012.

[19] Benjamin R. Handel, Igal Hendel, and Michael D. Whinston. Equilibria in health exchanges: Adverse selection vs. reclassification risk. 2012.

[20] Benjamin R. Handel and Jonathan Kolstad. Health insurance for "humans": Information frictions, plan choice, and consumer welfare. 2013.

[21] Amanda Kowalski. Estimating the tradeoff between risk protection and moral hazard with a nonlinear budget set model of health insurance. 2013.

[22] Thomas G. McGuire, Jacob Glazer, Joseph P. Newhouse, Sharon-Lise Normand, Julie Shi, Anna D. Sinaiko, and Samuel H. Zuvekas. Integrating risk adjustment and enrollee premiums in health plan payment. *Journal of Health Economics*, 32:1263–1277, 2013.

[23] Thomas G. McGuire, Joseph P. Newhouse, Sharon-Lise Normand, and Samuel H. Zuvekas. Assessing incentives for service-level selection in private health insurance exchanges. *Journal of Health Economics*, forthcoming, 2014.

[24] Joseph P. Newhouse, Mary Price, Jie Huange, J. Michael McWilliams, and John Hsu. Steps to reduce favorable risk selection in medicare advantage largely succeeded, boding well for health insurance exchanges. *Health Affairs*, 31:2618–2628, 2012.

[25] Maria Polyakova. Regulation of insurance with adverse selection and switching costs: Evidence from medicare part d. 2013.

[26] Michael Rothschild and Joseph Stiglitz. Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. *The Quarterly Journal of Economics*, 90:629–649, 1976.

[27] Julie Shi. Efficiency in plan choice with risk adjustment and premium discrimination in health insurance exchanges. 2013.

[28] Anna D. Sinaiko and Richard A. Hirth. Consumers, health insurance, and dominated choice. *Journal of Health Economics*, 30:450–457, 2011.

[29] Johannes Spinnewijn. Heterogeneity, demand for insurance and adverse selection. 2013.

[30] Kenneth Train. *Train, Kenneth*. Cambridge University Press, 2009.

[31] Jane M. Zhu, Timothy J. Layton, Anna D. Sinaiko, and Thomas G. McGuire. The power of reinsurance in health insurance exchanges to improve the fit of the payment system and reduce incentives for adverse selection. *Inquiry*, forthcoming, 2014.

Appendix 1: Risk Adjustment Model Estimation

In this paper I simulate risk adjustment with four different models: demographic, HHS-HCC, and prospective and concurrent industry models. In this appendix, I explain how the weights are estimated for each of these models. With model k, each individual is assigned a risk score r_i^k by multiplying a set of risk adjusters, X_i^k, by a set of risk adjustment weights, β^k:

$$r_i^k = X_i^k \beta^k$$

For the demographic model, X_i^k consists of a set of age-by-sex cells. For the other models, X_i^k consists of the same age-by-sex cells plus a set of Hierarchical Condition Categories (HCCs). HCCs are indicator variables equal to one if an individual receives a diagnosis that maps to the HCC.All diagnoses come from health insurance claims. Each of the over 14,000 ICD-9-CM diagnoses map to one of 394 industry HCCs and 264 HHS HCCs. For both the industry and HHS models, a subset of the HCCs are chosen for inclusion in the model based on explanatory power and concerns about efficiency. The models are also differentiated by the time period from which the diagnoses generating the HCCs are drawn. For the HHS-HCC and industry concurrent models, the diagnoses are from the current year, while for the prospective industry model, the diagnoses are from the prior year. This results in the HHS-HCC and industry concurrent models having much more explanatory power than the industry prospective model.

All models are estimated using data from the Truven Health Analytics Marketscan Commercial Claims and Encounters database. For some subset of the individuals in the dataset (usually those enrolled in a fee-for-service health plan), the total annual healthcare costs. c_i, are calculated and HCCs are assigned. The risk adjustment weights are then estimated via the following linear regression of individual costs on the risk adjusters:

$$c_i = X_i^k \delta^k + \varepsilon_i^k$$

The estimated weights are then normalized by dividing by the average cost in the sample:

$$\beta^k = \hat{\delta}^k / \bar{c}.$$

I note that all of the risk scores assigned in this paper use pre-estimated weights from software provided by either HHS or Verisk Health.

The estimation of the HHS-HCC weights varies slightly in the definition of c_i. Instead of using total costs, the HHS-HCC model is estimated using "plan liability" costs. Plan liability costs are defined by finding an individual's out-of-pocket costs based on their total costs and some standard non-linear cost sharing schedule found using the HHS actuarial value calculator. Because actuarial value varies by plan tier, plan liability costs also vary. This results in 5 separate models, one for each tier in the Exchanges. In the simulations, this implies that an individual's risk score depends on whether she enrolls in the Bronze or the Platinum plan, with risk scores generally being higher in the Platinum plan. Because of this, the HHS-HCC transfer formula includes an adjustment for the actuarial value of the plan relative to the enrollment-weighted average actuarial value in the market:

$$T_j = \left(\frac{\bar{R}_i^j}{\bar{R}} - \frac{AV^j}{\bar{AV}} \right) \bar{P}$$

I implement this adjustment in the simulations of the HHS-HCC model but not in the simulations of the industry models because with those models individuals' risk scores do not depend on the plan in which they choose to enroll.

Appendix 2: Out-of-Pocket Cost Distributions of New Enrollees

As explained in Section 4.2.3, an individual's expected out-of-pocket distribution, $F_{ij}(OOP)$, is estimated based on a predictive measure of the individual's future out-of-pocket costs, λ_{it}. This measure is calculated using sophisticated predictive modeling software from Verisk Health, and it incorporates virtually all of the information available from the individual's health insurance claims from the prior year. Unfortunately, for the new enrollees who were not enrolled in a plan during the prior year, the information necessary to calculate λ_{it} is not available, so a different method for estimating $F_{ij}(OOP)$ is necessary. I follow Handel and Kolstad (2013) by calculating $\lambda_{i,t+1}$ for each new enrollee using data from the current year. I then use the estimate of $\lambda_{i,t+1}$ to determine the probability that an individual with a predicted cost of $\lambda_{i,t+1}$was in each of the 500 cells at time t. I then determine the parameters of the individual's expected out-of-pocket cost distribution, $F_{ij}(OOP)$, by multiplying the individual's probability of being in a given cell in time t, given that the individual has predicted cost $\lambda_{i,t+1}$ in time $t+1$, by the estimated parameters for that cells. This produces a Bayesian estimate of the individual's expected out-of-pocket cost distribution, given $\lambda_{i,t+1}$.

Figure 1: Equilibrium sorting with adverse selection

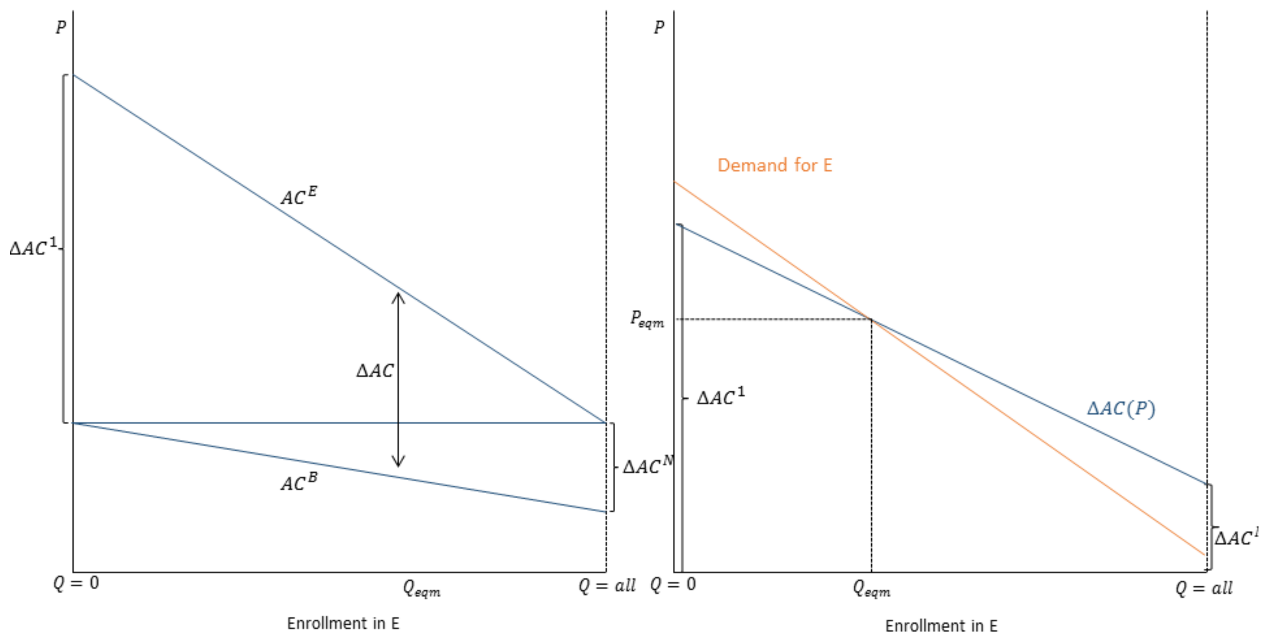

Notes: In both panels, enrollment in Plan E is on the x-axis, and the price differential is on the y-axis. The left panel describes the average cost of individuals enrolled in Plan E and Plan B as a function of enrollment in E. The difference between the lines defines the incremental average cost curve shown in the right panel. The right panel illustrates the equilibrium price differential and sorting where the demand curve crosses the incremental average cost curve. The figure illustrates the case where Plan E is adversely selected.

Figure 2: Adverse selection on risk scores

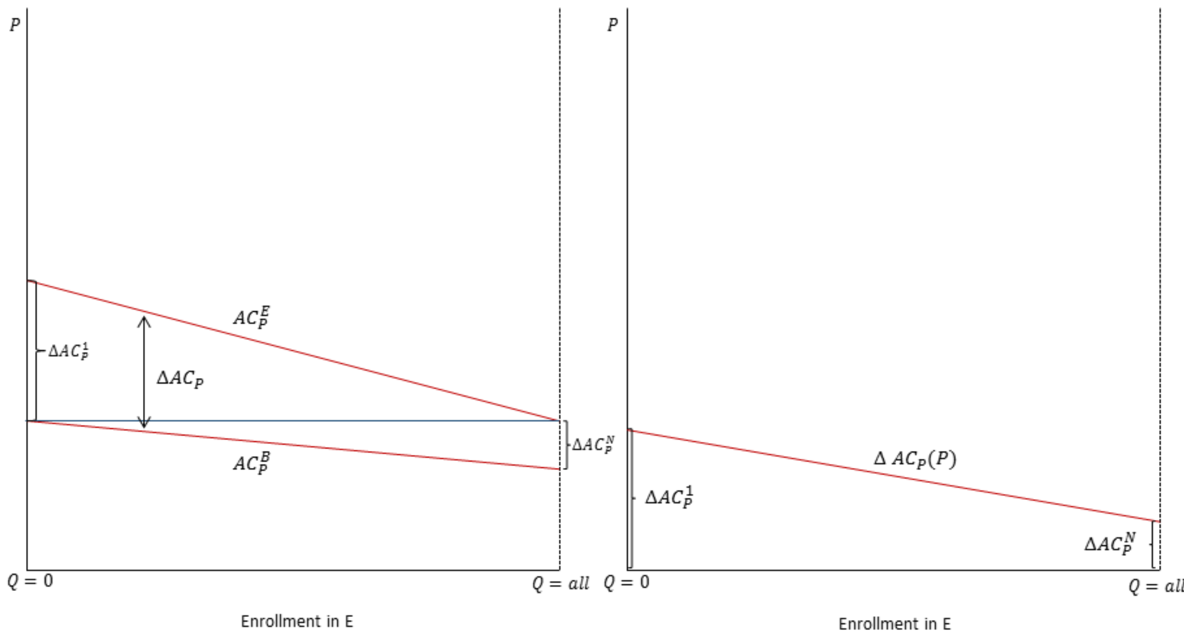

Notes: In both panels, enrollment in Plan E is on the x-axis, and the price differential is on the y-axis. The left panel describes the average predicted cost of individuals enrolled in Plan E and Plan B as a function of enrollment in E. The difference between the lines defines the incremental average predicted cost curve shown in the right panel. The figure illustrates the case where Plan E is adversely selected on predicted costs. Predicted costs are a function of r_i, the measure of risk upon which risk adjustment transfers are based.

Figure 3: Average Costs and Risk Adjustment Transfers

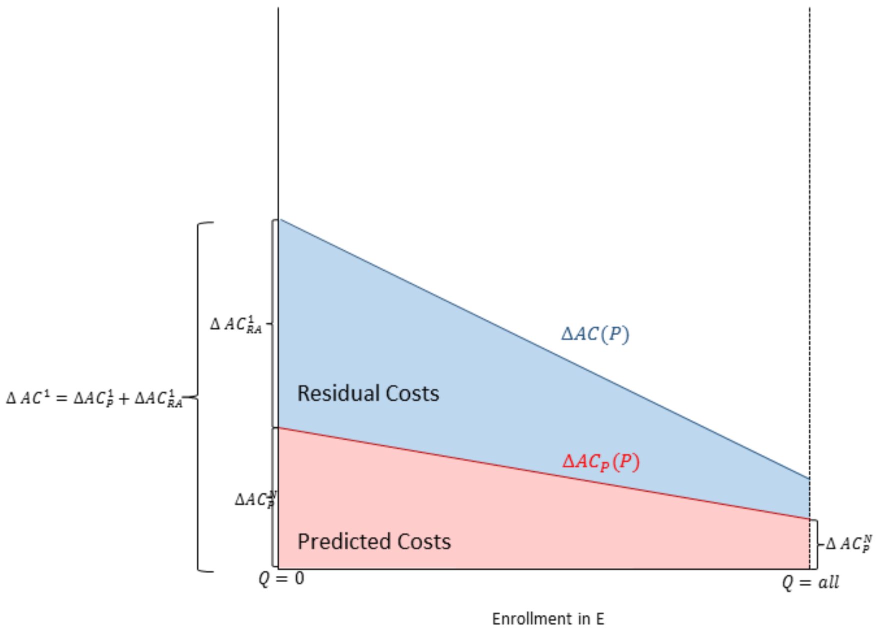

Notes: In the figure, enrollment is on the x-axis and the price differential is on the y-axis. The figure combines the left panels of Figures 1 and 2 to clearly separate total costs into predicted costs and residual costs. Predicted costs are the costs explained by the risk adjustment model, and residual costs are all other costs. Under risk adjustment, predicted costs are pooled across all plans in the market, thus eliminating that portion of the difference in costs across plans.

Figure 4: Equilibrium with Risk Adjustment

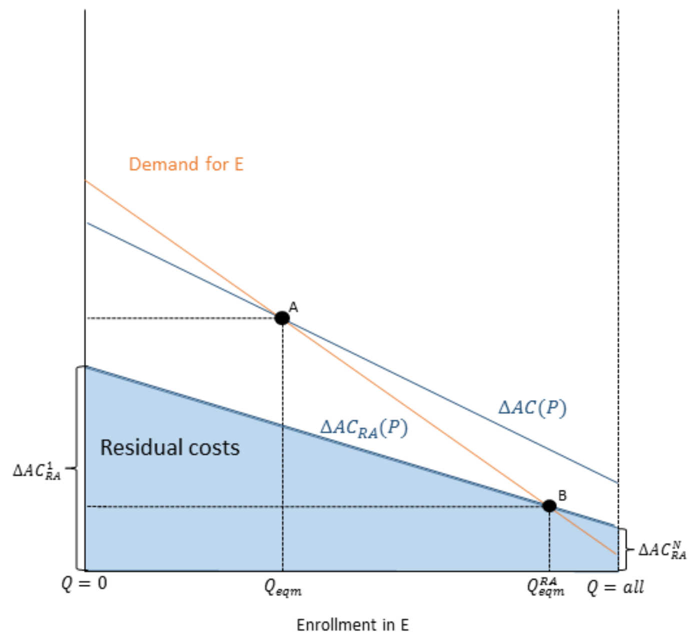

Notes: In the figure, enrollment is on the x-axis and the price differential is on the y-axis. The figure shows the incremental average cost curve with and without risk adjustment. Without risk adjustment the incremental average cost curve, $\Delta AC(P)$, is based on total costs. With risk adjustment the incremental average risk adjusted cost curve, $\Delta AC_{RA}(P)$, is based only on residual costs, or the portion of costs not explained by the risk adjustment model. Equilibrium without risk adjustment occurs at point A and equilibrium with risk adjustment occurs at point B. Enrollment in E is higher and the premium differential is lower with risk adjustment in this case.

Figure 5: Equilibrium with Risk Adjustment

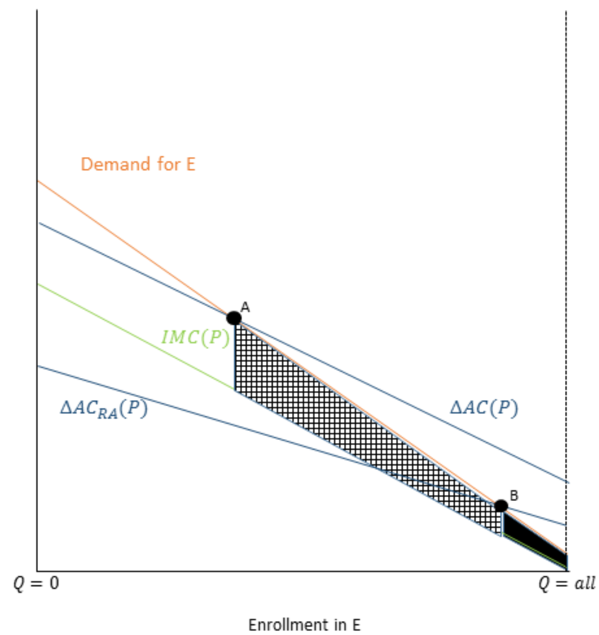

Notes: In the figure, enrollment is on the x-axis and the price differential is on the y-axis. The figure shows the incremental average cost curve with and without risk adjustment. Equilibrium without risk adjustment occurs at point A and equilibrium with risk adjustment occurs at point B. The green line describes the incremental marginal cost curve, which, along with demand, is assumed to be linear in the price differential. In this case, the welfare loss from adverse selection when there is no risk adjustment is described by the combination of the two shaded areas. The welfare loss with risk adjustment is described only by the solid shaded area and is substantially smaller in this case.

Figure 6: Correlation between Willingness-to-Pay for Platinum and Total Costs

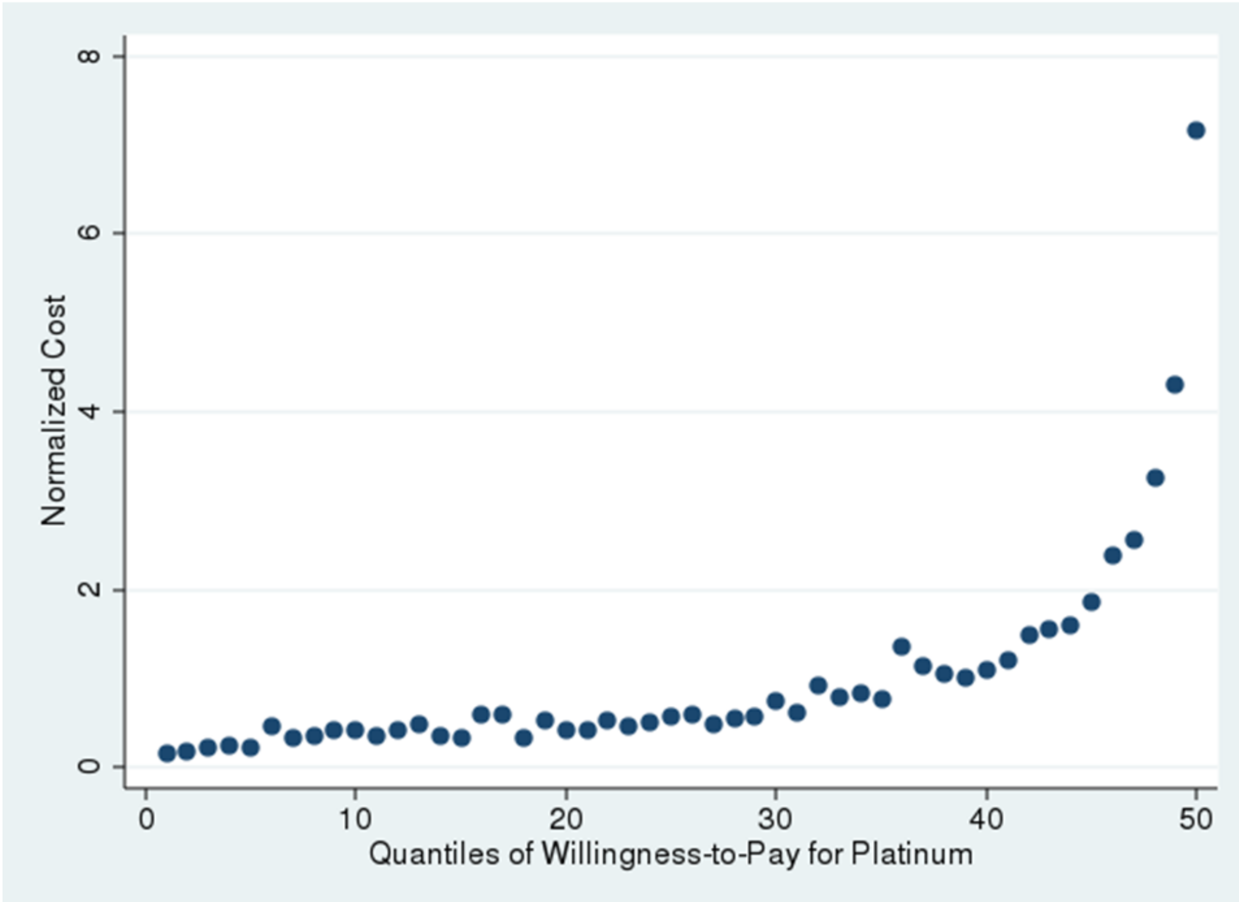

Notes: In the figure, 50 quantiles of willingness-to-pay are on the x-axis and normalized total medical spending (spending divided by average spending) is on the y-axis. Willingness-to-pay is generated using the expected utility model presented in the text along with the parameters of the model estimated using the choice model. A positive correlation between willingness-to-pay and total costs implies that as the price differential between Platinum and Bronze decreases, the marginal Platinum enrollee will have a lower cost than the average Platinum enrollee, i.e. the Platinum plan will be adversely selected.

Figure 7A: Total Cost and Demographic Risk

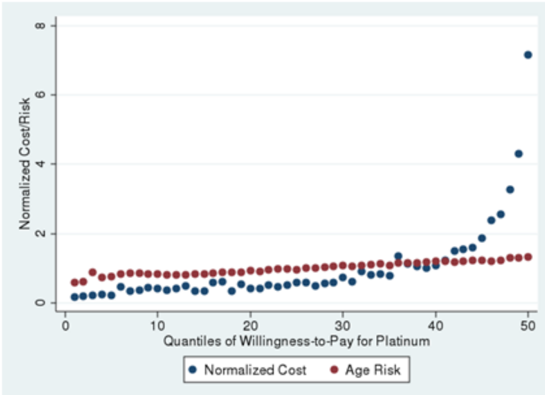

Figure 7B: Total Cost and Prospective Risk

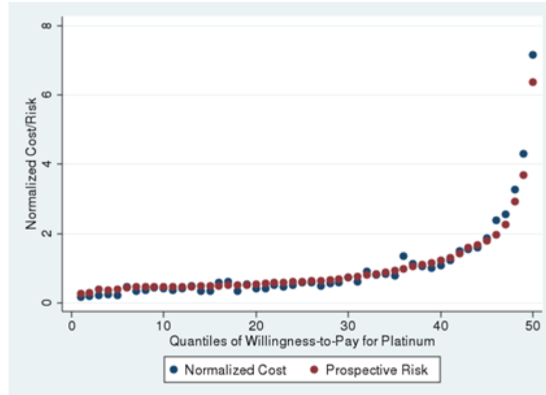

Figure 7C: Total Cost and Concurrent Risk

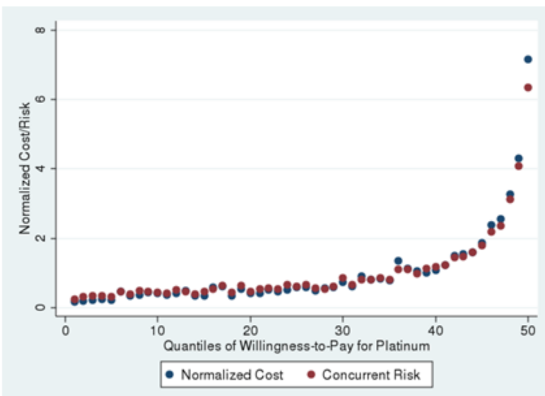

Figure 7C: Total Cost and HHS-HCC Risk

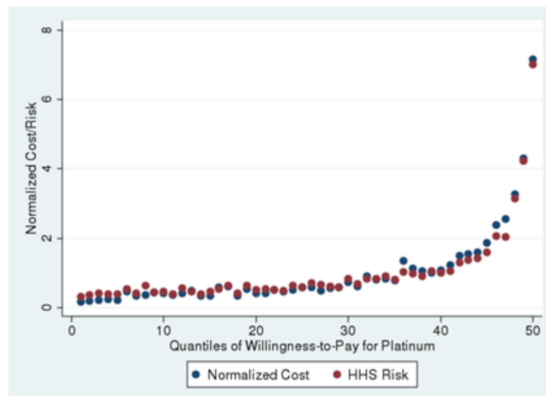

Notes: In the figures, 50 quantiles of willingness-to-pay are on the x-axis and normalized total medical spending (spending divided by average spending) and risk scores are on the y-axis. Willingness-to-pay is generated using the expected utility model presented in the text along with the parameters of the model estimated using the choice model. Risk scores are calculated using information from the health insurance claims data. This information is passed through software from Verisk Health and HHS to produce the risk scores. The software multiplies a vector of diagnostic- and demographic-based variables generated from the claims data by a vector of pre-estimated risk adjustment weights. A positive correlation between willingness-to-pay and total costs implies that as the price differential between Platinum and Bronze decreases, the marginal Platinum enrollee will have a lower cost than the average Platinum enrollee, i.e. the Platinum plan will be adversely selected. A positive correlation between willingness-to-pay and the risk score implies that as the price differential between Platinum and Bronze decreases, the marginal Platinum enrollee will have a lower risk score than the average Platinum enrollee, i.e. the Platinum plan will be adversely selected on predicted costs.

Figure 8: Equilibrium Search – Price Differential and Incremental Average Cost under Different Types of Risk Adjustment

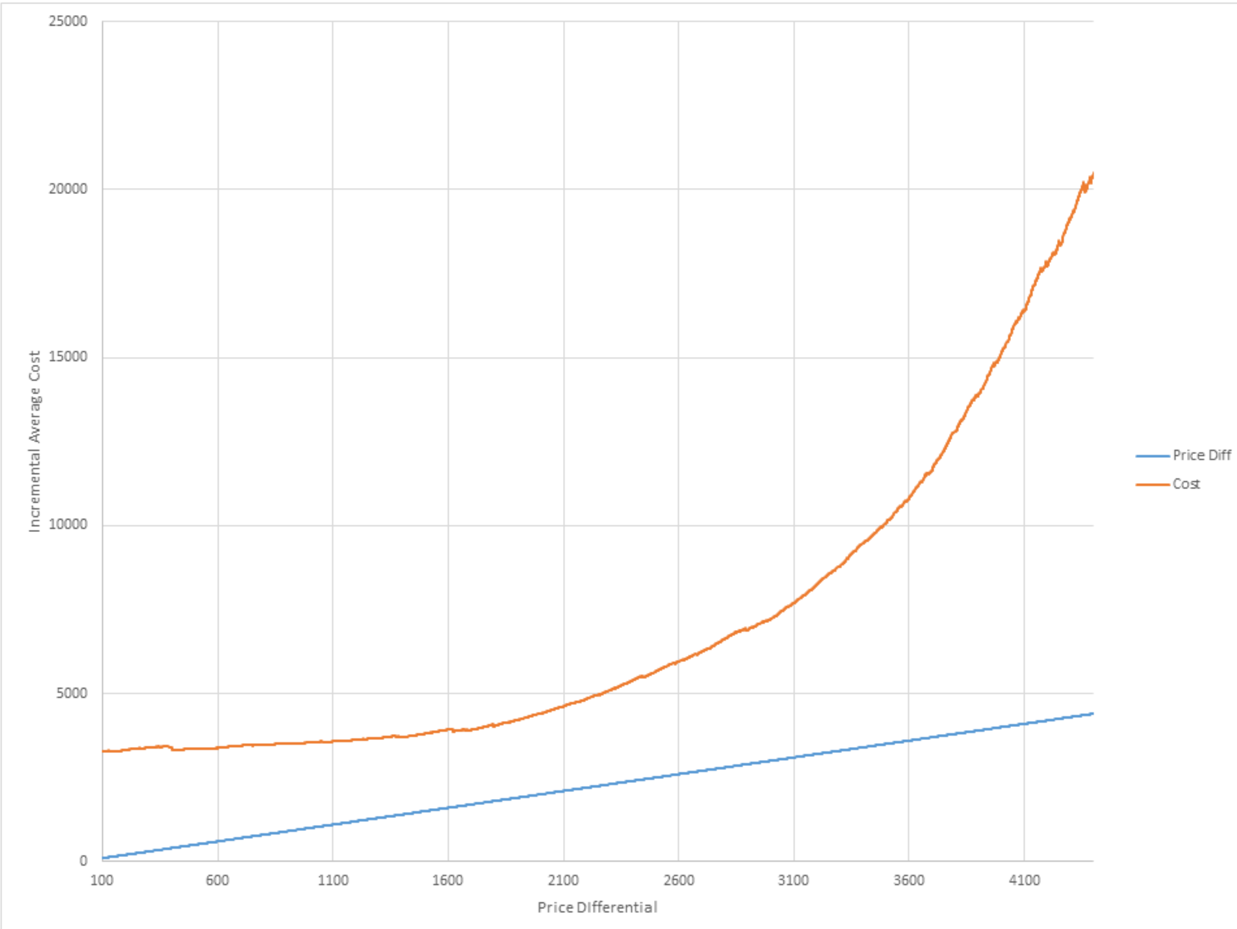

Notes: Figure shows search for equilibrium in setting where sample individuals required to choose between Bronze and Platinum Plans. Light blue line is the 45-degree line. Orange line represents incremental average cost (IAC) curve with no risk adjustment for each price differential. Because the IAC curve is everywhere above the 45-degree line, there is no equilibrium with positive enrollment in the Platinum plan, i.e. the market unravels. Prices, enrollment, and welfare can be found in Table 4.

Figure 9: Equilibrium Search

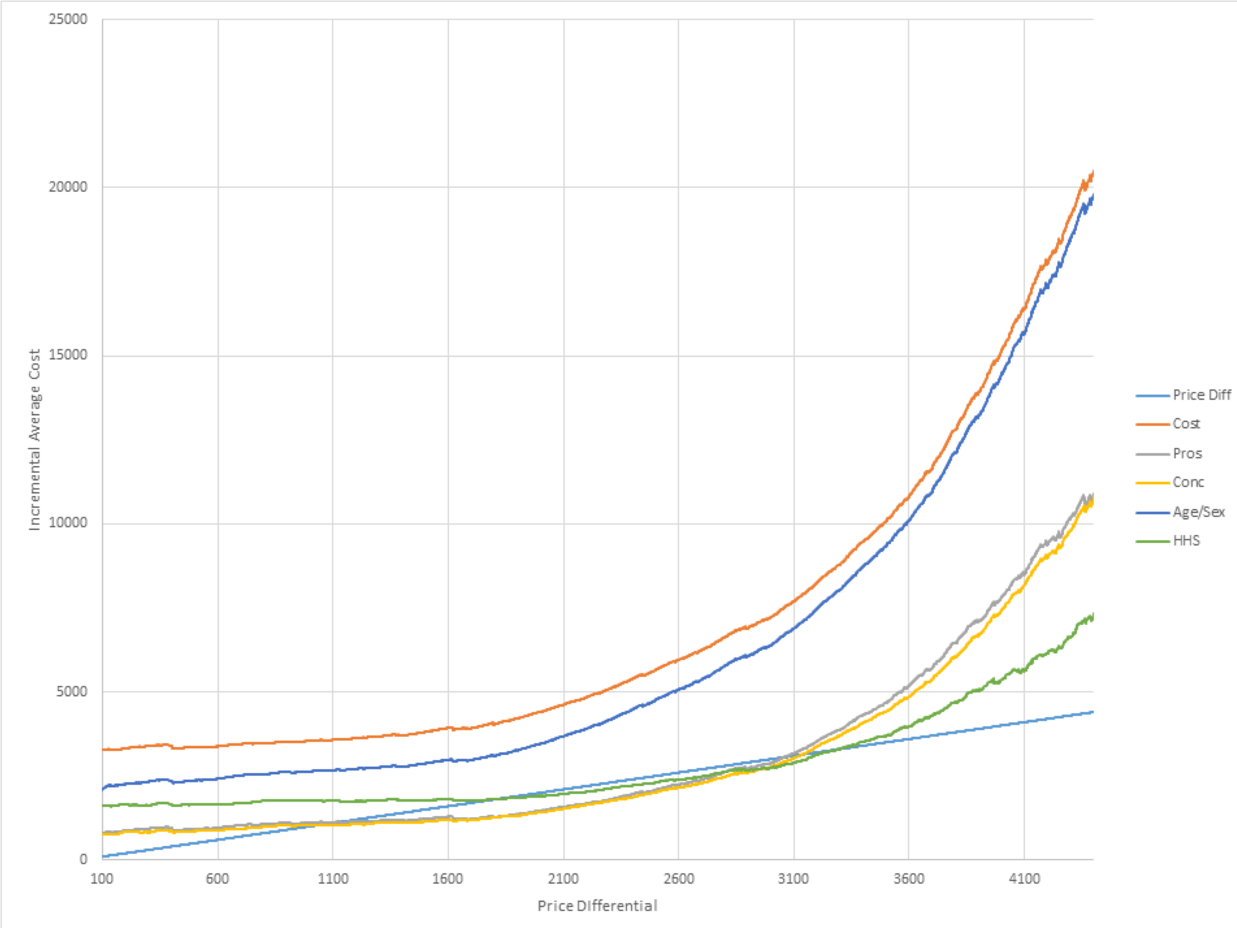

Notes: Figure shows search for equilibrium in setting where sample individuals required to choose between Bronze and Platinum Plans. Light blue line is the 45-degree line. Orange line represents incremental average cost (IAC) curve with no risk adjustment, blue line represents IAC with demographic risk adjustment, gray line represents IAC with prospective risk adjustment, gold line represents IAC with concurrent risk adjustment, and green line represents IAC with HHS-HCC risk adjustment. IAC with no and demographic risk adjustment is everywhere above 45-degree line implying complete market unraveling where everyone enrolls in Bronze plan. Prospective, concurrent, and HHS risk adjustment IACs cross 45-degree line, implying an interior equilibrium exists. Equilibrium is at lowest P where IAC crosses 45-degree line. Concurrent results in the lowest price differential. Prices, enrollment, and welfare can be found in Table 4.

Figure 10: Equilibrium Price Differentials and Sorting under Different Types of Risk Adjustment

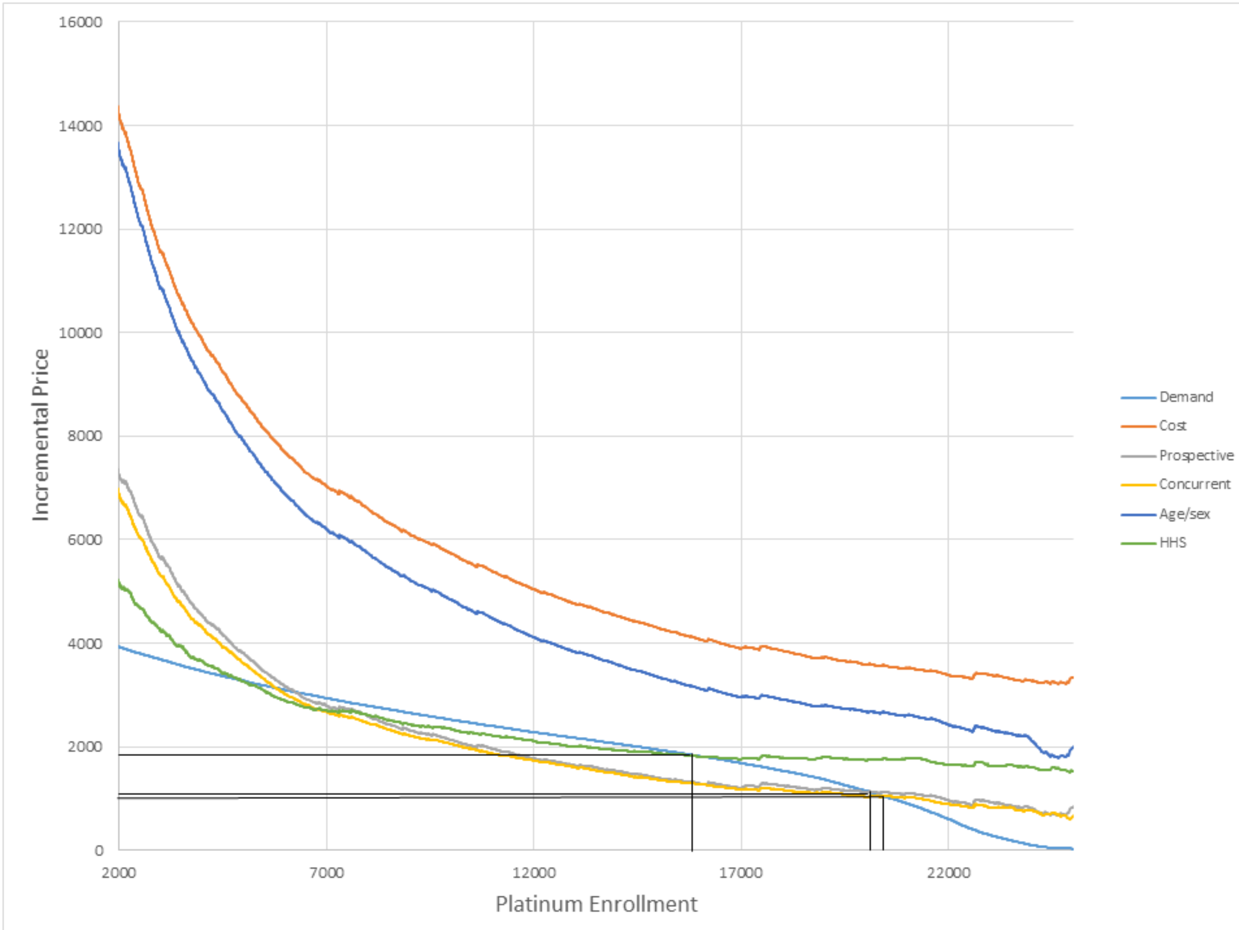

Notes: Figure shows equilibrium in setting where sample individuals required to choose between Bronze and Platinum Plans. Light blue line is the demand curve. Orange line represents incremental average cost (IAC) curve with no risk adjustment, blue line represents IAC with demographic risk adjustment, gray line represents IAC with prospective risk adjustment, gold line represents IAC with concurrent risk adjustment, and green line represents IAC with HHS-HCC risk adjustment. Enrollment in Platinum Plan is on x-axis. IAC with no and demographic risk adjustment is everywhere above 45-degree line implying complete market unraveling where everyone enrolls in Bronze plan. Prospective, concurrent, and HHS risk adjustment IACs cross 45-degree line, implying an interior equilibrium exists. Equilibrium is at lowest P where IAC crosses 45-degree line. Concurrent results in the lowest price differential. Equilibrium price and enrollment in Platinum are highlighted by dotted lines. Prices, enrollment, and welfare can be found in Table 4.

Figure 11: Equilibrium Price Differentials and Sorting under Different Types of Risk Adjustment

(with Reinsurance)

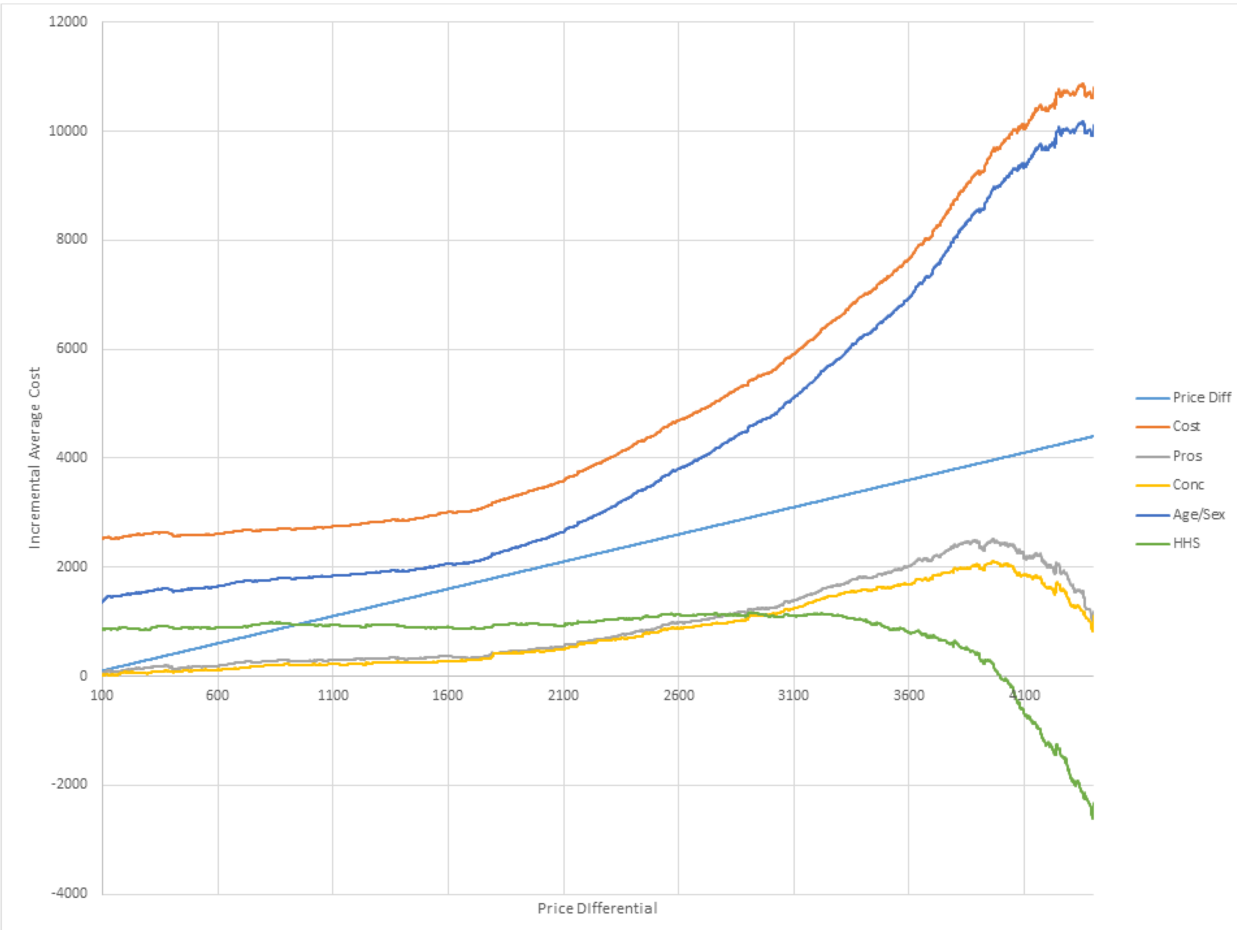

Notes: Figure shows search for equilibrium with reinsurance in setting where sample individuals required to choose between Bronze and Platinum Plans. Light blue line is the 45-degree line. Orange line represents incremental average cost (IAC) curve with no risk adjustment, blue line represents IAC with demographic risk adjustment, gray line represents IAC with prospective risk adjustment, gold line represents IAC with concurrent risk adjustment, and green line represents IAC with HHS-HCC risk adjustment. All lines include reinsurance. IAC with no and demographic risk adjustment is everywhere above 45-degree line implying complete market unraveling where everyone enrolls in Bronze plan. Prospective, concurrent, and HHS risk adjustment IACs cross 45-degree line, implying an interior equilibrium exists. Equilibrium is at lowest P where IAC crosses 45-degree line. Concurrent results in the lowest price differential. Prices, enrollment, and welfare can be found in Table 4.

Table 1: Summary Statistics

	Estimation Sample			Cost Model Sample	Simulation Sample
	Full sample	Old Employees	New Employees		
Male	0.6	0.6	0.67	0.49	0.6
New	0.08			n.a.	n.a.
Age:					
mean	41.16	41.61	35.95	34.98	41.5
1st Pcntl	21	22	20	1	22
25th Pcntl	32	32	26	17	32
50th Pcntl	45	42	34	38	42
75th Pcntl	50	50	46	51	50
99th Pcntl	62	62	61	63	62
Total costs:					
mean	3875.44	3998.5	2422.19	3086.05	3792.1
1st Pcntl	0	0	0	0	0
5th Pcntl	0	0	0	0	0
25th Pcntl	73	90	0	129	85
50th Pcntl	857	910	368	566.43	900
75th Pcntl	3016	3136	1811	2074.46	3109
95th Pcntl	15789	16167	9019	11970.83	15097
99th Pcntl	52385	53019	28672	40495	47342
Expected costs:					
mean	3959.34	4028.07	3147.69		3804.45
1st Pcntl	317.64	332.65	283.52		29.42
5th Pcntl	479.48	496.89	354.11		65.83
25th Pcntl	1073.26	1111.25	757.62		780.63
50th Pcntl	2071.85	2132.63	1461.67		1792.84
75th Pcntl	4119.02	4207.45	3041.46		3932.96
95th Pcntl	12201.57	12368.52	9711.7		12600.92
99th Pcntl	30161.83	30407.75	22785.8		33447.02
Total cost risk scores:					
mean	0.91	0.93	0.71	0.89	0.92
1st Pcntl	0.15	0.15	0.15	0.1	0.15
5th Pcntl	0.18	0.18	0.15	0.14	0.18
25th Pcntl	0.31	0.33	0.23	0.27	0.32
50th Pcntl	0.53	0.55	0.38	0.51	0.54
75th Pcntl	0.98	1	0.77	0.99	1.01
95th Pcntl	2.65	2.73	2.13	2.59	2.74
99th Pcntl	6.19	6.28	4.5	5.87	6.18
Prospective RA risk scores	0.89	0.9	0.75		0.89
Concurrent RA risk scores	0.83	0.85	0.58		0.81
N	9133	8420	713	845000	25398

Notes: Summary statistics for Estimation Sample, Cost Model Sample, and Simulation Sample. All samples come from Truven Marketscan dataset from choice years 2006-08. Estimation and Simulation Samples are from one large firm in Marketscan dataset where employees choose between 2 PPO plans. Samples are restricted to single-coverage employees enrolled for all 365 days of the year prior to and year of plan choice to ensure that costs can be predicted using full set of information. Estimation sample is restricted to employees from choice year 2007. Cost model sample is formed by first taking all individuals in Marketscan during at least 300 days of both 2006-2007. Then, total cost risk scores are generated from prior health claims using Verisk Health DxCG predictive modeling software. Marketscan sample is divided into 1000 cells based on total cost risk scores from year 1. Cost Model Sample is constructed by taking a random sample of 1000 individuals from each cell. Lognormal distribution with point mass at zero fit to costs in year 2 for each cell. Expected costs calculated by finding mean of the estimated distribution. Prospective and concurrent risk scores generated using Verisk Health DxCG risk adjustment software using diagnoses and demographics from year 1.

Table 2: Cost-sharing Parameters for Firm and Simulation Plans

	Firm		Simulations	
	Basic	Enhanced	Bronze	Platinum
Deductible	500	300	4500	$0
Coinsurance	20%	10%	20%	20%
OOP Max	4750	2600	6500	1500
Drug copay	Generic = $10 Brand=$5	Generic = $10 Brand=$5	Included in medical deductible, coinsurance, OOP max	

Notes: Table shows cost-sharing parameters for plan options at the firm and for plan options in the Exchange simulations. Firm parameters are used to create $f_{ij}(OOP)$ for estimation of the choice model, simulation parameters are used to create $f_{ij}(OOP)$ for Bronze and Platinum plans in the simulations. Under all plans, consumers pay the full cost of care up to the deductible, then they pay the coinsurance rate up to the out-of-pocket max. Beyond the out-of-pocket max, the consumer pays nothing. For the firm plans, drug coverage is not part of the price schedule, but coverage is identical in the two plans. For the simulations, drug spending is included with other medical spending in the non-linear price schedule. In the firm plans, ER visits and preventive visits are free of charge, but these visits make up only a small portion of total medical expenditures, so they considered to be priced with other medical spending.

Table 3: Choice Model Results

	Parameter Estimate	Parameter Std Error
Enhanced Shifter	-178.22	34.87
Switching Cost - Intercept	1585.74	407.39
Switching Cost - Age Coeff	21.91	12.84
Switching Cost - Fem Coeff	-359.24	258.70
CARA - Intercept	$1.0*10^{-3}$	$1.7*10^{-4}$
CARA - Log expend	$-3.2*10^{-5}$	$1.7*10^{-5}$
CARA - Age Coeff	$-8.4*10^{-6}$	$3.9*10^{-6}$
CARA - Age*Log expend	$4.2*10^{-7}$	$3.5*10^{-7}$
CARA - Fem Coeff	$2.3*10^{-5}$	$6.4*10^{-5}$
CARA - Pros Risk Coeff	$-3.1*10^{-5}$	$2.1*10^{-5}$
CARA - Conc Risk Coeff	$2.7*10^{-6}$	$1.5*10^{-5}$
Preference Shock - Std Dev	20.07	231.50
Mean CARA	$6.3*10^{-4}$	
Median CARA	$6.2*10^{-4}$	
Mean Switching Cost	2424.70	
Median Switching Cost	2418.24	

Notes: Results from simulated maximum likelihood estimation of choice model described in the paper. Enhanced shifter is a plan-specific intercept for the Enhanced Plan. Switching costs are estimated by comparing the choices of switchers and those of new enrollees of similar age and gender. CARA intercept represents the coefficient of absolute risk aversion for a zero year-old with total predicted cost and risk adjustment risk scores of zero. Log expend coefficient describes how CARA parameter varies with the log of total realized costs. Pros and conc risk coefficients describe how CARA parameter varies with risk adjustment risk scores. Mean CARA parameter implies that average individual in the sample would be indifferent between the status quo and a lottery that offered $100 with 50% probability and $95 with 50% probability.

Table 4: Equilibrium Prices, Sorting, and Welfare with Uniform Pricing

	No Reinsurance				
	Price Differential	Bronze Price	Platinum Price	% in Platinum	Change in Welfare
No Risk Adjustment	n.a.	$1,969	n.a.	0.0%	n.a.
Age/sex Risk Adjustment	n.a.	$1,969	n.a.	0.0%	$0
Prospective Risk Adjustment	$1,131	$2,357	$3,497	79.1%	$759
Concurrent Risk Adjustment	$1,050	$2,423	$3,473	80.4%	$767
HHS Risk Adjustment	$1,806	$1,992	$3,799	63.4%	$627
	Reinsurance				
	Price Differential	Bronze Price	Platinum Price	% in Platinum	Change in Welfare
No Risk Adjustment	n.a.	$1,969	n.a.	0%	n.a.
Age/sex Risk Adjustment	n.a.	$1,969	n.a.	0%	$0
Prospective Risk Adjustment	$35	$3,377	$3,374	98%	$798
Concurrent Risk Adjustment	n.a.	n.a.	$3,380	100%	$798
HHS Risk Adjustment	$946	$2,502	$3,449	82.0%	$775

Notes: Table shows equilibrium price differential (price of Platinum – price of Bronze), prices, proportion enrolled in Platinum plan, and change in welfare from no risk adjustment case to case with indicated type of risk adjustment. Bottom panel adds reinsurance where reinsurance reimburses 85% of an individual's plan costs above $60,000 and is funded with an actuarially fair per capita premium. Equilibrium found using algorithm described in the text. If there is no interior equilibrium, there is no price differential, and only the price of the plan in which the entire market enrolls is shown. Types of risk adjustment include age/sex which uses only demographic variables to predict costs, prospective and concurrent which use prior and current diagnosis groups, respectively, to predict costs, and HHS which is a concurrent model that uses a different set of diagnosis groups and allows for higher risk scores for Platinum enrollees and a penalty factor for the Platinum plan. Welfare calculated by the certainty equivalent concept discussed in the paper.

Table 5: Equilibrium Prices, Sorting, and Welfare with Age-based Pricing

	No Reinsurance				
	Avg Price Diff	Avg Bronze Price	Avg Platinum Price	% in Platinum	Change in Welfare
No Risk Adjustment	n.a.	$1,969	n.a.	0%	n.a.
Age/sex Risk Adjustment	n.a.	$1,969	n.a.	0%	$0
Prospective Risk Adjustment	$1,194	$2,291	$3,494	79%	$752
Concurrent Risk Adjustment	$1,121	$2,347	$3,475	80%	$760
HHS Risk Adjustment	$2,027	$1,850	$3,873	56%	$551
	Reinsurance				
	Avg Price Diff	Avg Bronze Price	Avg Platinum Price	% in Platinum	Change in Welfare
No Risk Adjustment	n.a.	$1,969	n.a.	0%	n.a.
Age/sex Risk Adjustment	n.a.	$1,969	n.a.	0%	$0
Prospective Risk Adjustment	$46	$3,301	$3,375	98%	$797
Concurrent Risk Adjustment	n.a.	n.a.	$3,380	100%	$798
HHS Risk Adjustment	**$1,030**	**$2,424**	**$3,456**	**81.9%**	**$769**

Notes: Table shows equilibrium price differential (price of Platinum – price of Bronze), prices, proportion enrolled in Platinum plan, and change in welfare from no risk adjustment case to case with indicated type of risk adjustment. Bottom panel adds reinsurance where reinsurance reimburses 85% of an individual's plan costs above $60,000 and is funded with an actuarially fair per capita premium. In all cases, prices vary by age according to the HHS age curve. Equilibrium found using algorithm described in the text. If there is no interior equilibrium, there is no price differential, and only the price of the plan in which the entire market enrolls is shown. Types of risk adjustment include age/sex which uses only demographic variables to predict costs, prospective and concurrent which use prior and current diagnosis groups, respectively, to predict costs, and HHS which is a concurrent model that uses a different set of diagnosis groups and allows for higher risk scores for Platinum enrollees and a penalty factor for the Platinum plan. Welfare calculated by the certainty equivalent concept discussed in the paper. Case in bold represents the full pricing policy currently being implemented in the Exchanges.

Table 6: Equilibrium Prices, Sorting, and Welfare with Uniform Pricing (1000 cells)

	No Reinsurance				
	Price Differential	Bronze Price	Platinum Price	% in Platinum	Change in Welfare
No Risk Adjustment	n.a.	$1,969	n.a.	0.0%	n.a.
Age/sex Risk Adjustment	n.a.	$1,969	n.a.	0.0%	$0
Prospective Risk Adjustment	$1,102	$2,394	$3,492	79.4%	$819
Concurrent Risk Adjustment	$1,006	$2,461	$3,468	81.0%	$826
HHS Risk Adjustment	$1,858	$1,972	$3,830	63.2%	$676
	Reinsurance				
	Price Differential	Bronze Price	Platinum Price	% in Platinum	Change in Welfare
No Risk Adjustment	n.a.	$1,969	n.a.	0%	n.a.
Age/sex Risk Adjustment	n.a.	$1,969	n.a.	0%	$0
Prospective Risk Adjustment	$66	$3,301	$3,367	95%	$856
Concurrent Risk Adjustment	n.a.	n.a.	$3,380	100%	$856
HHS Risk Adjustment	$929	$2,521	$3,450	82.2%	$832

Notes: Table shows equilibrium price differential (price of Platinum – price of Bronze), prices, proportion enrolled in Platinum plan, and change in welfare from no risk adjustment case to case with indicated type of risk adjustment. All simulations use choice model parameters and expected out-of-pocket cost distributions with 1,000 cells of λ_i. Bottom panel adds reinsurance where reinsurance reimburses 85% of an individual's plan costs above $60,000 and is funded with an actuarially fair per capita premium. In all cases, prices vary by age according to the HHS age curve. Equilibrium found using algorithm described in the text. If there is no interior equilibrium, there is no price differential, and only the price of the plan in which the entire market enrolls is shown. Types of risk adjustment include age/sex which uses only demographic variables to predict costs, prospective and concurrent which use prior and current diagnosis groups, respectively, to predict costs, and HHS which is a concurrent model that uses a different set of diagnosis groups and allows for higher risk scores for Platinum enrollees and a penalty factor for the Platinum plan. Welfare calculated by the certainty equivalent concept discussed in the paper. Case in bold represents the full pricing policy currently being implemented in the Exchanges.

Table 7: Equilibrium Prices, Sorting, and Welfare with Uniform Pricing (50 cells)

	No Reinsurance				
	Avg Price Diff	Avg Bronze Price	Avg Platinum Price	% in Platinum	Change in Welfare
No Risk Adjustment	n.a.	$1,969	n.a.	0%	
Age/sex Risk Adjustment	n.a.	$1,969	n.a.	0%	$0
Prospective Risk Adjustment	$1,114	$2,379	$3,489	79%	$708
Concurrent Risk Adjustment	$1,053	$2,420	$3,473	80%	$714
HHS Risk Adjustment	$2,042	$1,978	$4,025	52%	$468
	Reinsurance				
	Avg Price Diff	Avg Bronze Price	Avg Platinum Price	% in Platinum	Change in Welfare
No Risk Adjustment	n.a.	$1,969	n.a.	0%	
Age/sex Risk Adjustment	n.a.	$1,969	n.a.	0%	$0
Prospective Risk Adjustment	$68	$3,373	$3,376	99%	$749
Concurrent Risk Adjustment	n.a.	n.a.	$3,380	100%	$749
HHS Risk Adjustment	**$938**	**$2,507**	**$3,445**	**81.7%**	**$723**

Notes: Table shows equilibrium price differential (price of Platinum – price of Bronze), prices, proportion enrolled in Platinum plan, and change in welfare from no risk adjustment case to case with indicated type of risk adjustment. All simulations use choice model parameters and expected out-of-pocket cost distributions with 50 cells of λ_i. Bottom panel adds reinsurance where reinsurance reimburses 85% of an individual's plan costs above $60,000 and is funded with an actuarially fair per capita premium. In all cases, prices vary by age according to the HHS age curve. Equilibrium found using algorithm described in the text. If there is no interior equilibrium, there is no price differential, and only the price of the plan in which the entire market enrolls is shown. Types of risk adjustment include age/sex which uses only demographic variables to predict costs, prospective and concurrent which use prior and current diagnosis groups, respectively, to predict costs, and HHS which is a concurrent model that uses a different set of diagnosis groups and allows for higher risk scores for Platinum enrollees and a penalty factor for the Platinum plan. Welfare calculated by the certainty equivalent concept discussed in the paper. Case in bold represents the full pricing policy currently being implemented in the Exchanges.

Table 8: Equilibrium Prices, Sorting, and Welfare with Uniform Pricing (10% Private Info)

	No Reinsurance				
	Price Differential	Bronze Price	Platinum Price	% in Platinum	Change in Welfare
No Risk Adjustment	n.a.	$1,969	n.a.	0.0%	n.a.
Age/sex Risk Adjustment	n.a.	$1,969	n.a.	0.0%	$0
Prospective Risk Adjustment	n.a.	$1,969	n.a.	0.0%	$0
Concurrent Risk Adjustment	$850	$2,618	$3,468	84.8%	$794
HHS Risk Adjustment	n.a.	$1,969	n.a.	0.0%	$0
	Reinsurance				
	Price Differential	Bronze Price	Platinum Price	% in Platinum	Change in Welfare
No Risk Adjustment	n.a.	$1,969	n.a.	0%	n.a.
Age/sex Risk Adjustment	n.a.	$1,969	n.a.	0%	$0
Prospective Risk Adjustment	$566	$2,856	$3,422	88%	$803
Concurrent Risk Adjustment	n.a.	n.a.	$3,380	100%	$808
HHS Risk Adjustment	$988	$2,511	$3,499	83.0%	$787

Notes: Table shows equilibrium price differential (price of Platinum – price of Bronze), prices, proportion enrolled in Platinum plan, and change in welfare from no risk adjustment case to case with indicated type of risk adjustment. All simulations use choice model parameters and expected out-of-pocket cost distributions with 10% private information about future health care costs. Bottom panel adds reinsurance where reinsurance reimburses 85% of an individual's plan costs above $60,000 and is funded with an actuarially fair per capita premium. In all cases, prices vary by age according to the HHS age curve. Equilibrium found using algorithm described in the text. If there is no interior equilibrium, there is no price differential, and only the price of the plan in which the entire market enrolls is shown. Types of risk adjustment include age/sex which uses only demographic variables to predict costs, prospective and concurrent which use prior and current diagnosis groups, respectively, to predict costs, and HHS which is a concurrent model that uses a different set of diagnosis groups and allows for higher risk scores for Platinum enrollees and a penalty factor for the Platinum plan. Welfare calculated by the certainty equivalent concept discussed in the paper. Case in bold represents the full pricing policy currently being implemented in the Exchanges.

Table 9: Equilibrium Prices, Sorting, and Welfare with Uniform Pricing (20% Private Info)

			No Reinsurance		
	Avg Price Diff	Avg Bronze Price	Avg Platinum Price	% in Platinum	Change in Welfare
No Risk Adjustment	n.a.	$1,969	n.a.	0.0%	$0
Age/sex Risk Adjustment	n.a.	$1,969	n.a.	0.0%	$0
Prospective Risk Adjustment	n.a.	$1,969	n.a.	0.0%	$0
Concurrent Risk Adjustment	$768	$2,693	$3,461	86%	$736
HHS Risk Adjustment	n.a.	$1,969	n.a.	0.0%	$0
			Reinsurance		
	Avg Price Diff	Avg Bronze Price	Avg Platinum Price	% in Platinum	Change in Welfare
No Risk Adjustment	n.a.	$1,969	n.a.	0.0%	
Age/sex Risk Adjustment	n.a.	$1,969	n.a.	0.0%	$0
Prospective Risk Adjustment	$599	$2,833	$3,433	88%	$741
Concurrent Risk Adjustment	n.a.	n.a.	$3,380	100%	$748
HHS Risk Adjustment	**$985**	**$2,526**	**$3,511**	**82.3%**	**$724**

Notes: Table shows equilibrium price differential (price of Platinum – price of Bronze), prices, proportion enrolled in Platinum plan, and change in welfare from no risk adjustment case to case with indicated type of risk adjustment. All simulations use choice model parameters and expected out-of-pocket cost distributions with 20% private information about future health care costs. Bottom panel adds reinsurance where reinsurance reimburses 85% of an individual's plan costs above $60,000 and is funded with an actuarially fair per capita premium. In all cases, prices vary by age according to the HHS age curve. Equilibrium found using algorithm described in the text. If there is no interior equilibrium, there is no price differential, and only the price of the plan in which the entire market enrolls is shown. Types of risk adjustment include age/sex which uses only demographic variables to predict costs, prospective and concurrent which use prior and current diagnosis groups, respectively, to predict costs, and HHS which is a concurrent model that uses a different set of diagnosis groups and allows for higher risk scores for Platinum enrollees and a penalty factor for the Platinum plan. Welfare calculated by the certainty equivalent concept discussed in the paper. Case in bold represents the full pricing policy currently being implemented in the Exchanges.

Appendix Figures

Figure A1: Equilibrium Search with Reinsurance

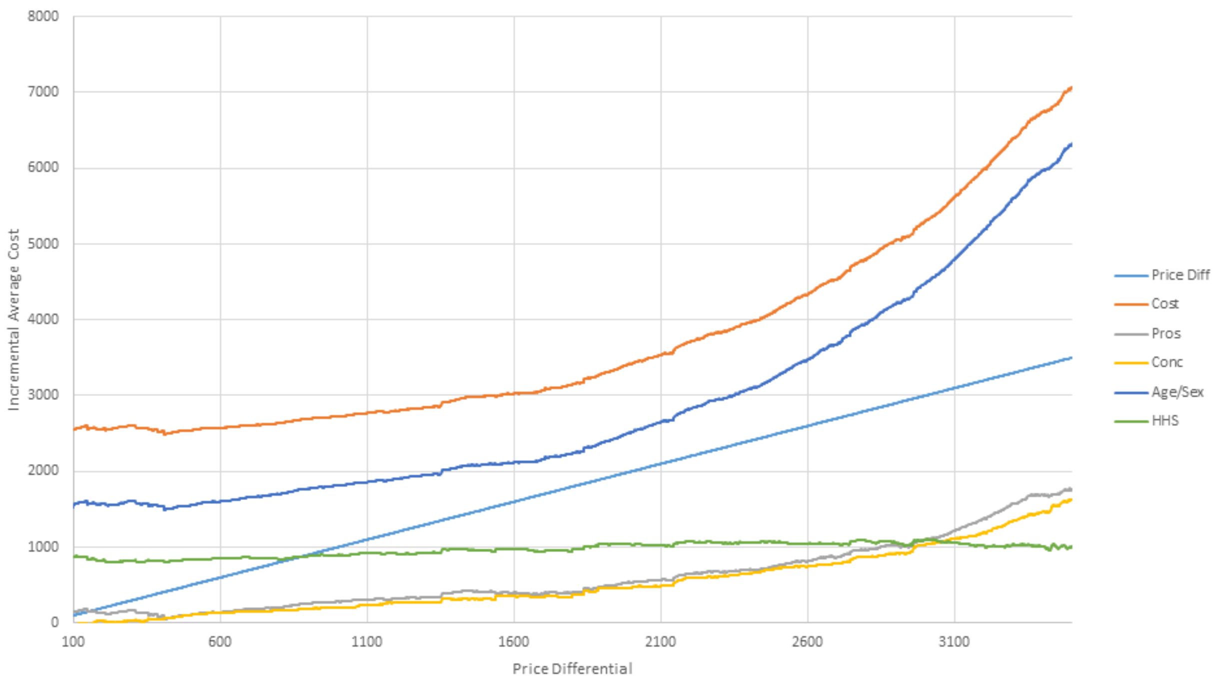

Figure A2: Equilibrium Search with Age-based Pricing

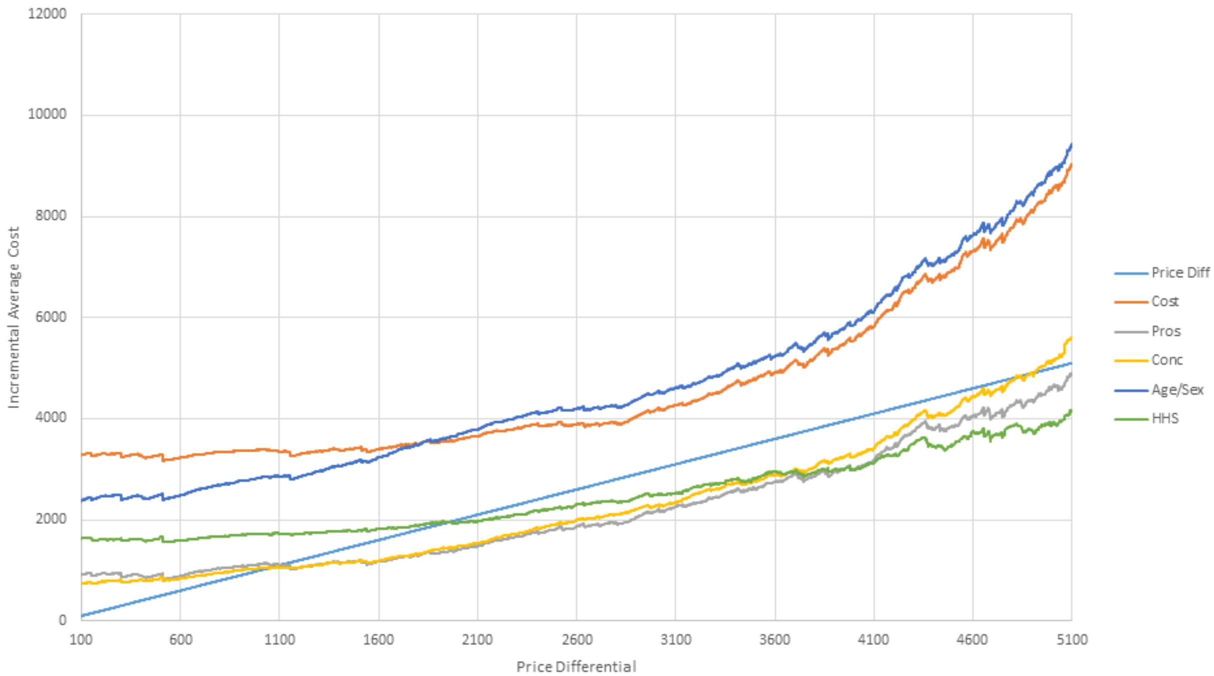

Figure A3: Equilibrium Search with Reinsurance and Age-based Pricing

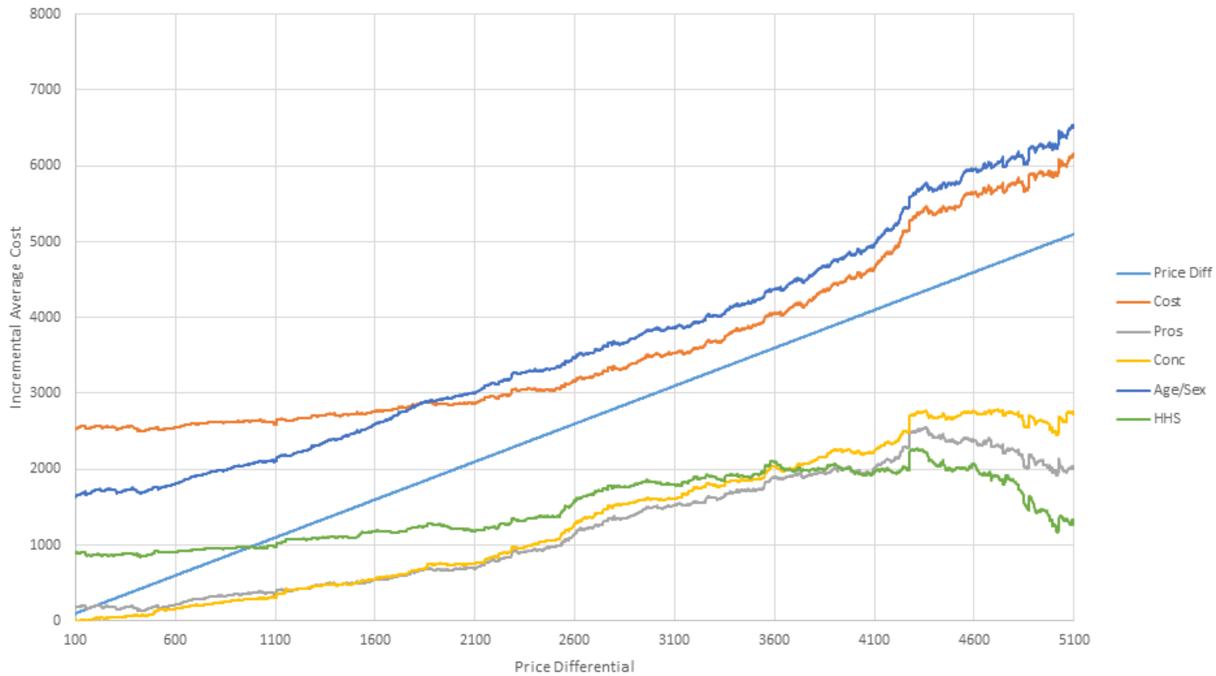

Figure A4: Advantageous selection on risk scores

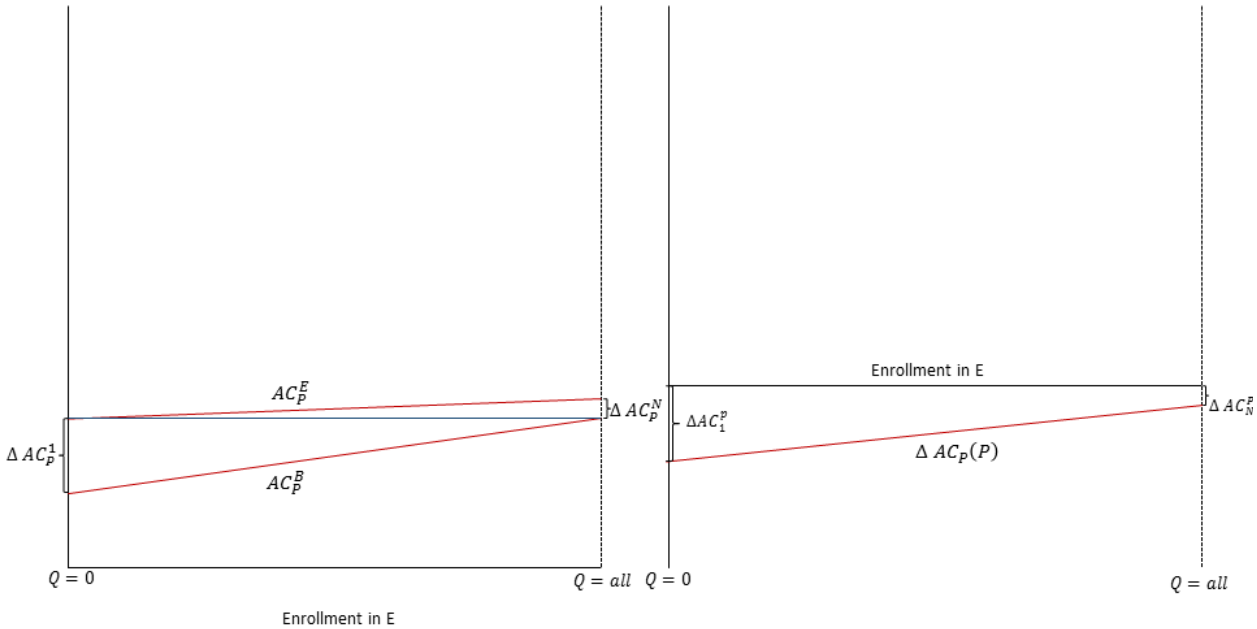

Notes: In both panels, enrollment in Plan E is on the x-axis, and the price differential is on the y-axis. The left panel describes the average predicted cost of individuals enrolled in Plan E and Plan B as a function of enrollment in E. The difference between the lines defines the incremental average predicted cost curve shown in the right panel. The figure illustrates the case where Plan E is advantageously selected on predicted costs. Predicted costs are a function of r_i, the measure of risk upon which risk adjustment transfers are based.

Figure A5: Equilibrium sorting with adverse selection

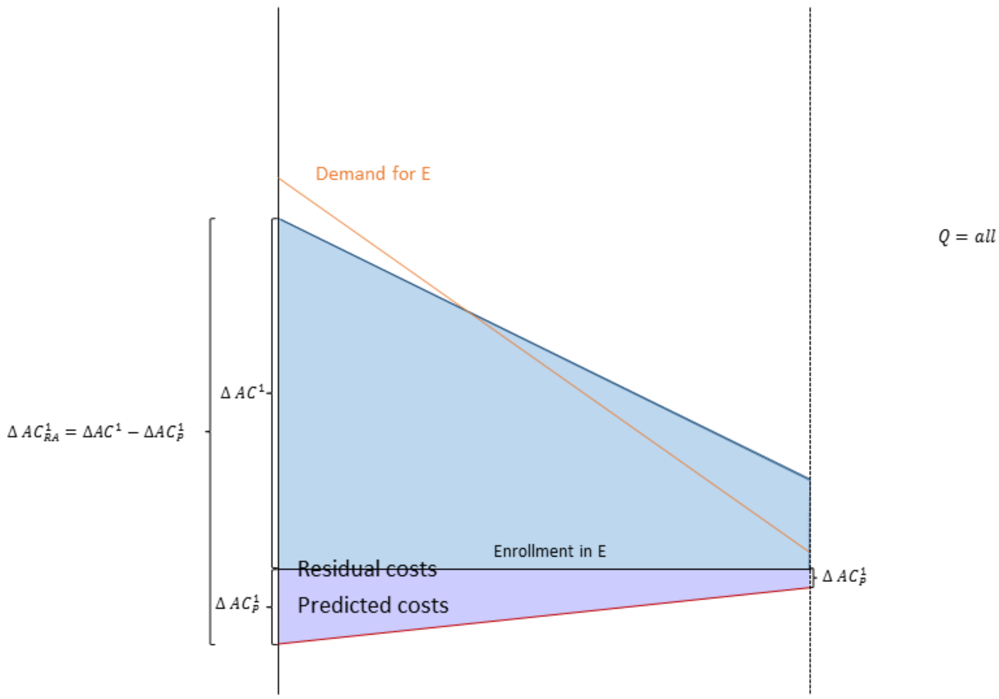

Notes: In the figure, enrollment is on the x-axis and the price differential is on the y-axis. The figure combines the left panels of Figures 1 and 2 to clearly separate total costs into predicted costs and residual costs. Predicted costs are the costs explained by the risk adjustment model, and residual costs are all other costs. Under risk adjustment, predicted costs are pooled across all plans in the market, thus eliminating that portion of the difference in costs across plans. In this case, predicted costs are negative, implying a negative transfer for Plan E and a positive transfer for Plan B.

Figure A6: Equilibrium sorting with adverse selection

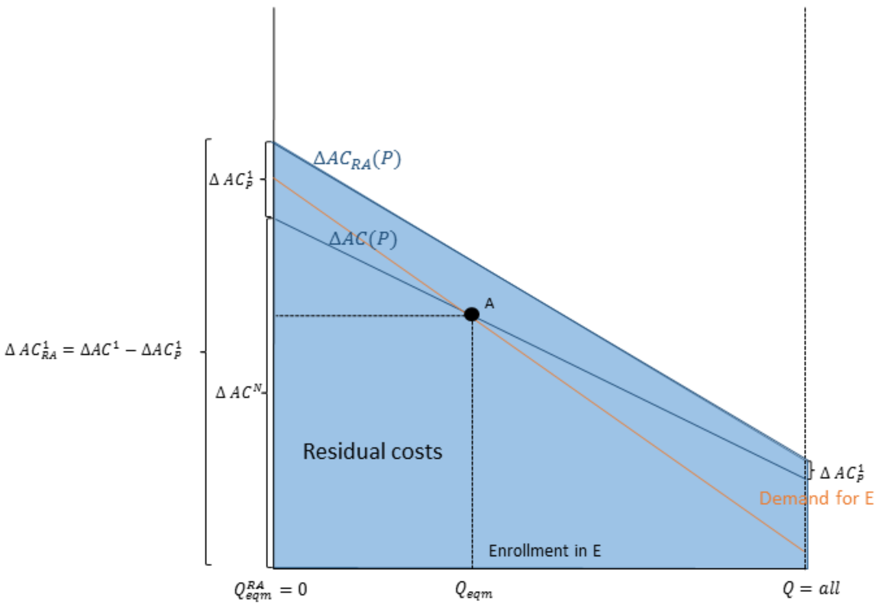

Notes: In the figure, enrollment is on the x-axis and the price differential is on the y-axis. The figure shows the incremental average cost curve with and without risk adjustment. Without risk adjustment the incremental average cost curve, $\Delta AC(P)$, is based on total costs. With risk adjustment the incremental average risk adjusted cost curve, $\Delta AC_{RA}(P)$, is based only on residual costs, or the portion of costs not explained by the risk adjustment model. Equilibrium without risk adjustment occurs at point A and equilibrium with risk adjustment occurs at the left edge of the figure where enrollment in E is equal to zero. Enrollment in E is lower and the premium differential is higher with risk adjustment in this case.