Horizon Effects and Adverse Selection in Health Insurance Markets

Olivier Darmouni and Dan Zeltzer*

February 14, 2017

Abstract

We study how increasing contract length affects adverse selection in health insurance markets. While health risks are persistent, private health insurance contracts in the U.S. have short, one-year terms. Short-term, community-rated contracts allow patients to increase their coverage only after risks materialize, leading to market unraveling. Longer contracts ameliorate adverse selection because both demand and supply exhibit horizon effects. Intuitively, longer horizon risk is less predictable, thus elevating demand for coverage and lowering equilibrium premiums. We estimate risk dynamics using data from 3.5 million U.S. health insurance claims and calculate counterfactual coverage and welfare in equilibrium with two-year contracts. We predict that such contracts would increase coverage by 6% from its initial level and yield average annual welfare gains of $100–$200 per person. Welfare gains from increased enrollment would partly offset by exposing those with low coverage to greater risk.

*Olivier Darmouni (omd2109@columbia.edu) is Assistant Professor of Finance and Economics at Columbia Business School, Columbia University, New York City, New York. Dan Zeltzer (dzeltzer@tauex.tau.ac.il) is Lecturer (Assistant Professor) of Economics at the School of Economics, Tel Aviv University, Tel Aviv, Israel. This paper was written while both authors were doctoral candidates at the Department of Economics at Princeton University. We thank Janet Currie and Sylvain Chassang for their advice. We also thank seminar participants at the Department of Economics and the Center for Health and Wellbeing at Princeton University. This research received financial support from the Program for U.S. Healthcare Policy Research of the Center for Health and Wellbeing at Princeton University. Access to the MarketScan Research Databases was provided by Truven Health Analytics graduate dissertation support program.

1 Introduction

While most U.S. health insurance contracts last one year, some illnesses last longer. The combination of predictable individual risk with frequent opportunity to adjust coverage may result in adverse selection. Mitigating selection is central to health insurance regulation, and it often involves limiting individual choice such as the enrollment mandate imposed by the Affordable Care Act (ACA). But with rising premiums on the one hand and enrollment in the marketplaces being concentrated almost exclusively in plans whose out-of-pocket exposure is high on the other hand, selection on both the extensive and intensive margins leaves many with incomplete risk protection.[1]

This study examines how extending the horizon of health insurance contracts impacts adverse selection. The main contribution is to explore the benefits of pooling risk over time *within* individuals, as opposed to simply *across* individuals. We show that increasing contract horizons can be used as an instrument to limit market unraveling. Conceptually, longer contracts reduce selection because individual risk is less predictable at longer horizons. This mechanism impacts both supply and demand, and it generally leads to greater coverage and improved welfare. We study these effects empirically, using administrative claims data on healthcare utilization. We estimate that risk predictability declines by 10%–15% per year due to mean-reversion in healthcare expenditures. Combining these risk estimates with a model of insurance markets, we predict that increasing contact horizons from one to two years would increase coverage by 6% and reduce deadweight loss on average by \$100–\$200.

We first show how to analyze the issue of risk dynamics and contract length using the framework of Einav et al. (2010). We focus on competitive community-rated insurance markets (like the ACA marketplaces), in which individuals periodically choose between standardized coverage levels. Graphically, increasing contract length induces a flattening of both cost and demand curves. Intuitively, because individual risk is mean-reverting, individual ability to predict future risk declines, making the risk pools more homogeneous. We show that supply response is unambiguous, while demand leans towards more high coverage if the mean type gets such coverage in the first place, and towards less coverage otherwise. In equilibrium, these combined effects generally lead to more coverage. Gains from improved coverage are partly offset by exposing those with low coverage to greater risk.

We then estimate the key parameter—risk predictability—over different horizons, and find that it declines over time in a significant and robust way. We use administrative data with detailed individual healthcare costs and utilization for two separate popula-

[1] In 2014, only 14% of enrollees chooses one of the two high coverage plans, Gold or Platinum. By 2016, high coverage dropped to 10% (Kaiser Foundation). On a related note, premiums on these plans have been rising over the years.

tions. Data from MarketScan Research Databases cover a working-age population with employer-sponsored insurance, and data from Medicare cover an elderly population with public insurance. We predict individual overall expenditure over different horizons using different sets of predictors. Risk predictability declines significantly within each year: we show that two-years-ahead risk is only 85% as predictable as one year ahead. Surprisingly, while predictability increases when more comprehensive predictors are used, its *decline* over time remains similar. This decline is also similar for the two different populations studied. This estimated decline in predictability reflects mean-reversion in risk. A complementary parameter—the persistence of health shocks—is shown, unsurprisingly, to be higher for the elderly population.

We use these risk-predictability estimates to evaluate the impact of increasing contract length on coverage and welfare. Given the lack of naturally occurring variation in contract length in the U.S., we estimate the demand response by calibrating a standard model of insurance choice in the line of Handel et al. (2015). We predict that such change would increase coverage by 6% over its initial level, yielding a reduction of 5-15% in deadweight loss (an average of 100–200 USD per person, depending on the initial level of coverage). Welfare gains are relatively small as longer contracts are a mixed blessing: Those who upgrade to high coverage plans enjoy welfare gains, but those who still remain with lower coverage, even when contracts are extended, suffer from longer exposure to risk without the ability to upgrade coverage.

This study adds to the literature studying adverse selection in health insurance markets and potential policy responses. Previous works have suggested the presence of such selection (Cutler and Zeckhauser, 1998; Culter and Reber, 1998; Einav and Finkelstein, 2010; Heiss et al., 2013); studied its welfare implications (Einav et al., 2010; Hackmann et al., 2012); and considered policy responses, including: plan standardization (Rice et al., 1997), risk adjustment mechanisms (Ellis et al., 2000; McGuire et al., 2013; Brown et al., 2014; Bundorf et al., 2012; Geruso, 2013; Hackmann et al., 2015), and participation mandates Kolstad and Kowalski (2016). Bundorf et al. (2012) document preference heterogeneity which implies that preserving some choice among plans is useful, and show that adverse selection might not be dealt with by setting optimal premiums alone. Their findings highlight how our suggested policy instrument of mitigating adverse selection through increased contract duration is particularly useful, as it preserves choice while mitigating selection.

Only a few works however consider the dynamics of risk in health insurance contracts, including Aron-Dine et al. (2012), Handel et al. (2015), and Cabral (2016). Few theoretical works exist on adverse selection dynamics in health insurance, in settings radically different from the current regulatory environment, such as Diamond (1992) or Cochrane (1995). This paper, in contrast, studies an environment similar to the current ACA marketplaces. Handel et al. (2016) is the closest to this paper. They also explore the dynamics of risk and consider the impact of changing contract length on equilibrium coverage and welfare. The key difference is that they focus on contracts with very long (lifetime) horizons and assume that these contracts can be risk-rated. On the other hand, we focus on a more realistic reform of the existing marketplaces by looking at two-year, community-rated contracts.

2 The Mechanism

We model the equilibrium of a competitive health insurance exchange to show how increasing contract length affects adverse selection and coverage due to a decline in risk predictability over a longer horizon.

2.1 Setup

To illustrate the mechanism, consider a standard model of plan choice with adverse selection, that resembles the markets introduced by the ACA. Individuals face a risk process $\{c_t\}_{t>0}$ and know θ_0, an individual-specific parameter informative about these future costs. Individuals can choose one of two plans, High and Low. Denote by p the extra premium charged for the High plan over the default Low plan and by $\Delta\iota$ the difference in co-insurance rates. Individuals can only choose one plan per period. Define the *contract length* to be the length of the period between two insurance plan choices.

2.1.1 Market structure

Individuals purchase insurance contracts before the beginning of the contract period. We compare two cases with different contract lengths. In one case, the insurance contract length is one year. This baseline case—annual choice of contracts during an open enrollment period—is currently the most common setup in the U.S. It is in place in both the Affordable Care Act health insurance exchanges and most employer-sponsored health insurance markets.[2] In the other case, insurance contract length is two years. The key difference between one-year and two-year contracts is that with one-year contracts individuals may adjust their second-year coverage levels based on their first-year risk realization, whereas with two-year contracts they may not.This paper focuses on studying how increasing the contract length would affect coverage and welfare in equilibrium.

[2] Certain qualifying events allow some people to buy contracts for even shorter periods (see Aron-Dine et al., 2012).

In principle, one can study contracts longer than two years and even lifetime contracts (see Cochrane, 1995; Handel et al., 2016). We choose to focus on an incremental reform instead for two reasons: (i) feasibility, and (ii) confidence in extrapolating from annual contracts data.

We assume the market is a regulated health insurance exchange: contracts are community rated and rejections are prohibited, an individual mandate forces all individuals to participate, and insurers must offer a standardized menu of plans meeting certain actuarial values. For simplicity, we consider contracts are fully characterized by premium and coinsurance. Insurers are competitive.

2.1.2 Risk and Predictability

Consider a finite horizon model where individuals live for two years, $t = 1, 2$, facing a risk process of the following log-linear form:[3]

$$c_1 = \mu + \alpha_1 \theta_0 + \varepsilon_1 \tag{1a}$$

$$c_2 = \mu + \alpha_2 \theta_0 + \beta_1 \varepsilon_1 + \varepsilon_2 \tag{1b}$$

where c and θ_0 are the individual log cost and risk type (i.e., $C_t(\theta_0) = e^{1+c_t(\theta_0)}$ for $t = 1, 2$). The risk type θ_0 may include past expenditures as well as other characteristics. We assume θ_0, ε_1 and ε_2 are mutually independent, zero mean. That is, the risk of types $\theta_0 < 0$ is lower-than-average, and the risk of types with $\theta_0 > 0$ is higher-than-average. For simplicity of exposition, we also assume θ_0 is uniformly distributed on $[-\theta_{max}, \theta_{max}]$. We later relax this assumption and use administrative data to estimate this distribution.[4]

Horizon effects: The key parameters driving the effect of contract length are the persistence parameters α_1, α_2. They represent how predictive current risk type θ_0 is of future risk, respectively next year and two years from now. This is the channel by which contract length affects the distribution of risks in the population, and among other things, equilibrium coverage.

A summary statistic for equilibrium is the effective level of risk predictability, denoted by $\tilde{\alpha}$. The parameter $\tilde{\alpha}$ is defined as α_1 for one-period contracts, and $\frac{\alpha_1+\alpha_2}{2}$ for two-period contracts. Intuitively, risk is less predictable in the most distant future, so that $\alpha_1 > \alpha_2$. Therefore, effective risk predictability $\tilde{\alpha}$ falls with the horizon, as we empirically document in Section 3. In this way, changing the contract length is equivalent to doing comparative

[3]The model can easily be generalized to an infinite horizon AR model, or even ARIMA. Empirically, we use a log-linear specification that fits well the large skewness in health spending.

[4]The model does not explicitly account for moral hazard. See Section 3.2 for a discussion of the empirical implications of moral hazard in this setting.

statics with respect to the *effective risk predictability* $\tilde{\alpha}$. We use this approach to graphically represent the effect of increasing contract length on coverage and welfare.

2.1.3 Demand

For expositional purposes, assume a linear demand curve. Specifically, assume that (annualized) individual net utility from choosing the High insurance plan takes the form:

$$v = -p + \tilde{\xi}_\theta \theta_0 + \tilde{\xi}_0 \tag{2}$$

where tilde objects vary with contract length.[5] That is, we allow utility to flexibly vary with contract length but assume price sensitivity does not.[6]

Individuals purchase the High plan if their risk type θ_0 is above a cutoff, given by the marginal type θ^*:

$$\theta^*(p) = \frac{p - \tilde{\xi}_0}{\tilde{\xi}_\theta} \tag{3}$$

There is adverse selection: for a given price differential across plans, only risk types above a certain cutoff choose the High plan. The marginal type depends on price through the utility weight on private risk type $\tilde{\xi}_\theta$. In other words, $\tilde{\xi}_\theta$ governs demand elasticity, and, as we will explain below, it is the main channel by which contract length affects equilibrium via the demand side. As contract length increases, current risk type θ_0 is less predictive of future risk and therefore adverse selection is reduced.

2.1.4 Supply

The marginal cost curve represents the differential cost of providing higher coverage:

$$\Delta MC(p) = \Delta\iota(\mu + \tilde{\alpha}\theta^*(p)) \tag{4}$$

In competitive equilibrium, the premium difference between High and Low coverage reflects this difference in average cost between the risk pools. The average cost curve is:

$$\Delta AC(p) = \Delta\iota(\mu + \tilde{\alpha}\frac{\theta^*(p) + \theta_{max}}{2}) + \iota_l \tilde{\alpha}\theta_{max} \tag{5}$$

In equilibrium $\Delta AC(p) = p$, which leads to a price higher than the efficient price obtained

[5] Because v is a utility function, we normalize the coefficient on p to be -1 without loss of generality.

[6] In the two-year contract case, it is implicitly understood that premia are annualized so that p reflects per-year cost of choosing the High plan.

when $\Delta MC(p) = p$.[7] Intuitively, adverse selection leads to inefficiently low coverage in equilibrium. The appendix also develops a tractable closed-form solution of the case of CARA utility.

2.2 The Effects of Contract Length on Coverage

In this section, we perform the key comparative static exercise of the paper, namely what happens to coverage in equilibrium when contract length increases from one year to two years? The total effect is a sum of a supply channel and a demand channel.

2.2.1 The Supply Channel

Figure 1: Supply Response with Extended Contract Horizon

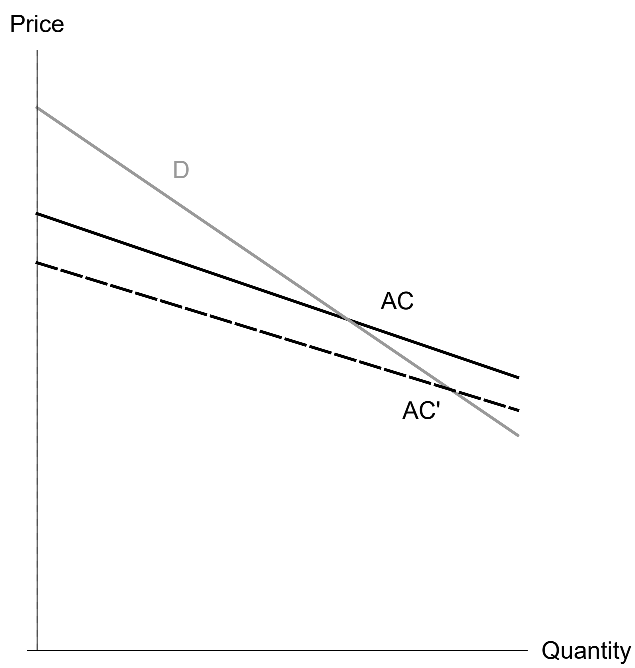

The first main result is that, on the supply side, an increase in contract length increases coverage. Intuitively, coverage increases through a flattening of the AC curve. Since AC shifts uniformly towards zero, this result is very general (Figure 1 illustrates this point).[8] As contract length increases, the effective risk predictability $\tilde{\alpha}$ falls because of mean reversion, which leads to a flattening of the marginal cost curve and to a decline in the difference

[7]The equilibrium price is given by $p^* = \frac{1}{1-\frac{\alpha\Delta\iota}{2\tilde{\xi}_\theta}} \left(\Delta\iota \left(\mu + \tilde{\alpha}\frac{\theta_{max} - \frac{\tilde{\xi}_\theta}{\tilde{\xi}_\theta}}{2} \right) + \iota_l \tilde{\alpha}\theta_{max} \right)$.

[8]All algebraic proofs are omitted, as they follow from elementary manipulations of the equilibrium price equation.

between risk pools for all coverage levels (Figure A1). From the point of an insurer, individuals become more homogeneous. For any given risk pool choosing the High plan at price p, the average risk falls. The difference between high- and low-risk types is more moderate, and consequently the dependence of insurers' profits on the composition of the risk pool is reduced. Graphically, the AC curve pivots counterclockwise toward zero. Under very general conditions, this change in the average cost curve induces an increase in equilibrium coverage.[9]

Increasing contract length reduces average cost by pooling risk over time, *within* an individual. This contrast with existing contracts that only pool risk *across* individuals. Longer contracts exploit the dynamics of risk: because risk displays some degree of mean-reversion, pooling risk over time reduces the sensitivity of insurers' profits to individual risk types.

2.2.2 The Demand Channel

Interestingly, the demand channel can increase or decrease coverage.

The intuition behind this result is illustrated in Figure 2. In both cases, an increase in contract length flattens the demand curve through a reduction in $\tilde{\xi}_\theta$. Significantly, this flattening occurs by pivoting around a point that has an important economic interpretation. The demand curve pivots counterclockwise around the valuation of the *representative* type $\bar{\theta}$.

Indeed, mean-reversion in health risk implies that type become effectively more homogeneous with longer contracts. The representative type $\bar{\theta}$ is defined as the type whose value of insurance is left unchanged by extending contract length. In the linear case, it corresponds to the average type $\bar{\theta} = 0$, as drawn in Figure 2.

This explains why the demand channel is ambiguous. Effectively, as types become more homogeneous, individual plan choices tend to mirror the plan choice of the representative type. In the left panel, the representative type originally does not choose the High plan. Therefore, as the contract length increases, fewer individuals choose the High plan. On the other hand, the right panel presents the case in which the representative type initially chooses the High plan. In this case, increasing contract length increases the share of the population choosing the High plan.

[9] Specifically, as long as the demand curve crosses the AC curve from above.

Figure 2: Demand Response and Equilibrium with Extended Contract Horizon

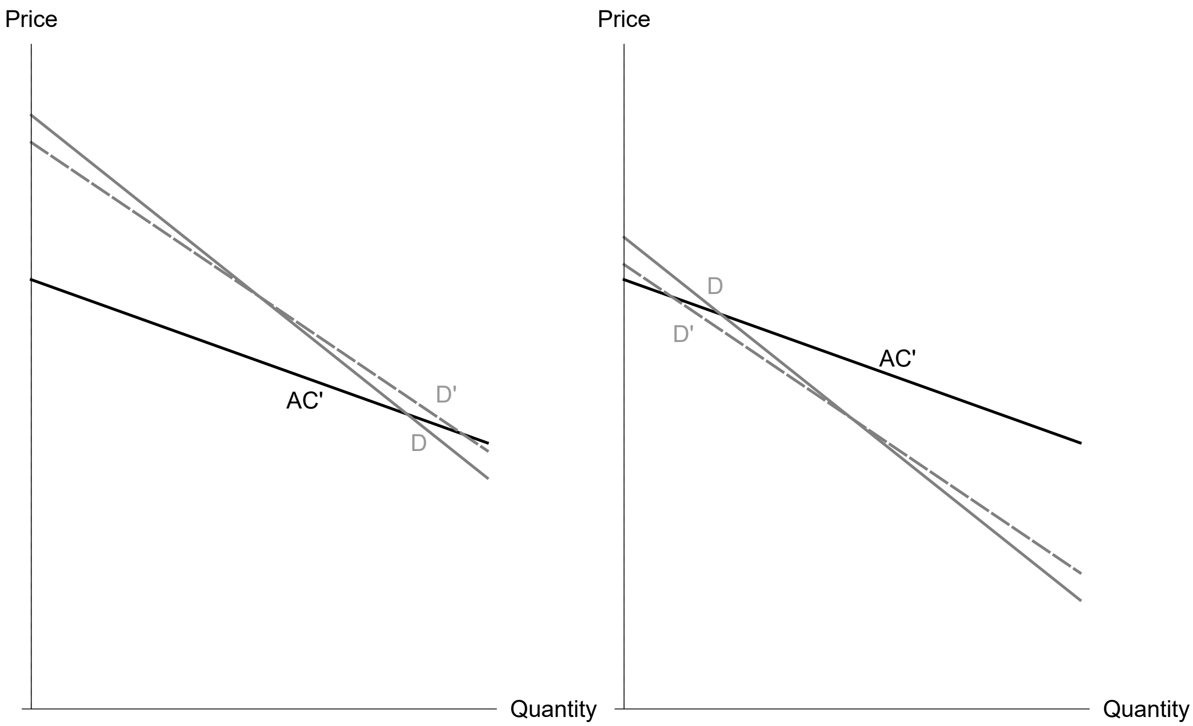

(a) $\theta^* > \bar{\theta}$ (b) $\theta^* < \bar{\theta}$

Notes: Equilibrium with one-year contracts (solid) and two-year contracts (dashed). The vertical dotted line denotes the quantity associated with $\theta = \bar{\theta}$. In case (a) the representative type purchases High coverage, the demand response augments the supply response and further increases equilibrium coverage. In case (b) the representative type does not purchase High coverage, the demand response dampens the supply response. In this example, coverage still increases in equilibrium, but less than in case (a).

2.2.3 Welfare

Figure 3 illustrates the welfare trade-off. For a given contract length, the welfare loss from adverse selection can be represented as a parallelogram. As contract length increases, there are two effects that go in opposite directions. On the one hand, more individuals choose the High plan, reducing the length of the welfare loss. On the other hand, the gap between the demand and the MC curve increases, increasing the height of the parallelogram. Intuitively, this second effect comes from the fact that individuals are exposed to more risk when contracts are longer because they are facing an additional interim shock ε_1, which (partly) persists until the next year.

Figure 3: Inefficiency Reduction due to Extended Contract Horizon

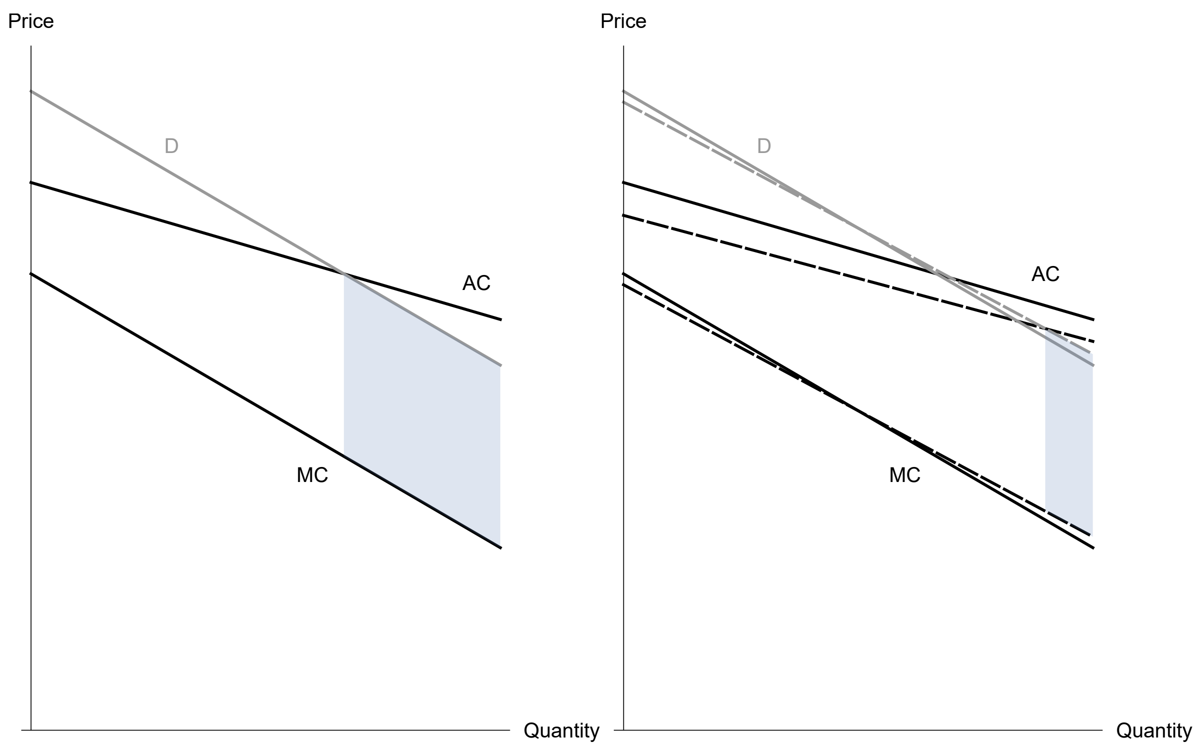

(a) one-year (b) two-year

Notes: Equilibria with one-year contracts (solid curves) and two-year contracts (dashed curves). The shaded area is the deadweight loss due to adverse selection. Extended contract horizon leads to better within-individual pooling, which reduces selection and the deadweight loss associated with it.

3 Empirical Estimates of Risk Predictability

We estimate the predictability of risk at different horizons, the key parameters that determine equilibrium coverage change. We estimate these parameters using healthcare utilization and cost data from two different administrative data sources: MarketScan Research Databases and Medicare Cost and Utilization Databases. We find similar results across these different samples: when moving from a one-year to a two-years horizon, risk predictability diminishes significantly.

3.1 Data

We use administrative data on healthcare utilization and costs. Important for our purpose, these data track individuals over time and, sourced from payors, are comprehensive. Data of the same granularity are used by insurers for planning and pricing, and thus the internal validity of our estimates is presumably high. Using two distinct sources covering different populations and different periods also supports their external validity. The rest of this section describes these datasets.

The first data source is Truven Health Analytics MarketScan Research Databases. These data capture individual clinical utilization, expenditures, and enrollment across inpatient, outpatient, prescription drug, and carve-out services from approximately 45 large employers for 2002–2004. These data represent the medical expenditures of the working-age population, an age profile similar to that of the target population of the ACA marketplaces.

The second data source is Medicare, the federal health insurance program for people who are 65 or older.[10] The data contain claims from Fee-For-Service Medicare beneficiaries. We use the Research Identifiable Files of the Master Beneficiary File and Cost and Utilization databases. These longitudinal databases record, for a 5% sample of Medicare beneficiaries, all healthcare utilization and costs over the period 2008–2012. For comparability with the population in the ACA marketplaces and the MarketScan sample, both of which cover working-age individuals, we restrict the sample to beneficiaries aged 65–70 when entering the sample.[11]

Table A3 shows summary statistics for each of the samples. The difference in age profiles and utilization patterns work in our favor, as they help us gauge the robustness of our estimates of the decline in risk predictability.

Figure 4 visualizes the key parameter that drives our results: mean reversion in risk (i.e., in healthcare costs). It is a binned scatter plot of annual log costs of our sample of 3.5 million MarketScan working-age enrollees. While risk is persistent—bins preserve their rank order year-over-year—it also exhibits substantial reversion to the mean. Our estimates

[10]We consider only elderly beneficiaries, although Medicare also covers certain younger people with disabilities, and people with end-stage renal disease

[11]For a complete list of utilization and cost measures used in the prediction of risks, see The Master Beneficiary Summary File and Cost and Use segment, http://www.resdac.org/cms-data/files/mbsf/data-documentation, accessed June 2016.

Figure 4: Mean Reversion in Individual Healthcare Costs

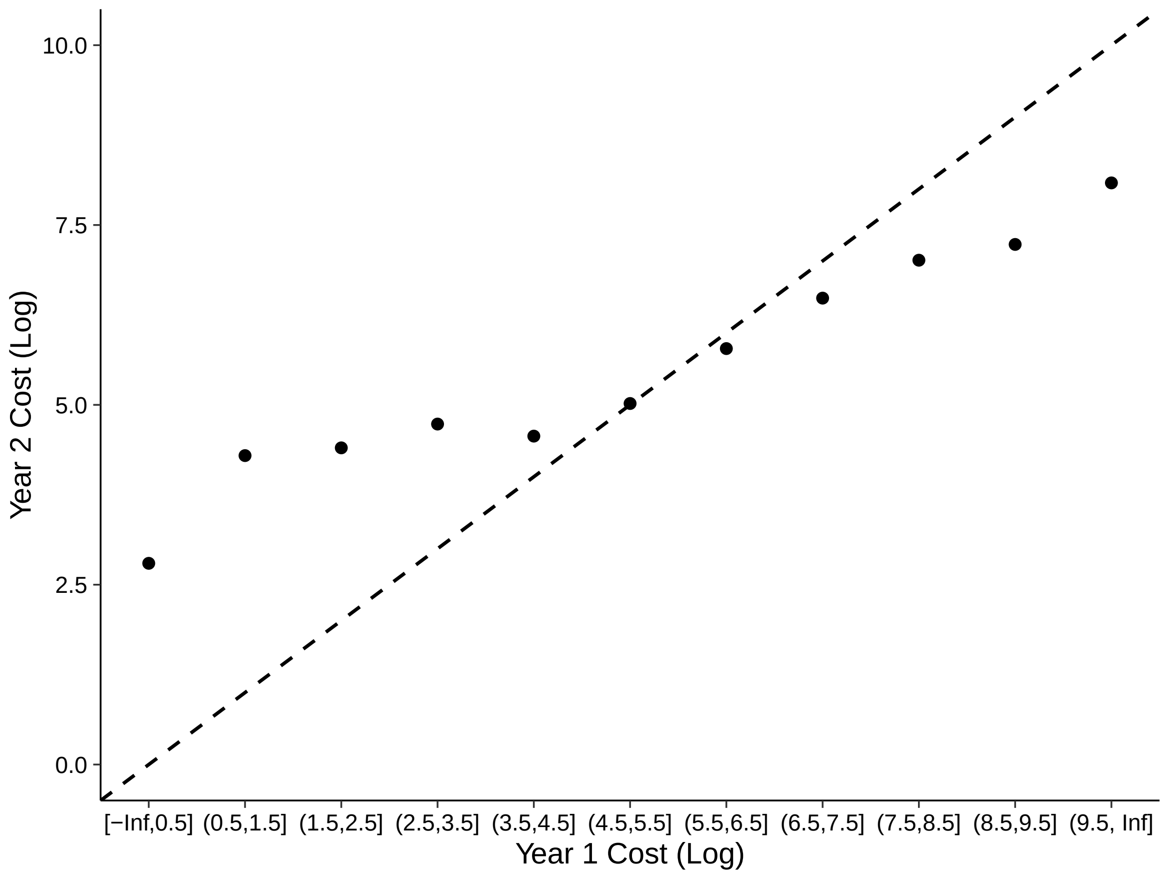

Notes: Binned scatter plot of individual healthcare costs over two subsequent years. Costs were aggregated at the individual level from administrative claims of 3.5 million working-age employees and their dependents from MarketScan Research Databases. The first-year costs were then binned into 11 equally spaced bins, with each dot representing the mean cost of the subsequent year for the individuals in the bin. The dashed line is the 45-degree line, representing unchanged costs over time. Differences between the means of each pair of bins are all significant ($p<0.01$).

of (1), presented below, show this reversion holds even when other predictors of risk are included. Figure A2 shows the sample distribution of log cost c_1.

3.2 Empirical Model of Risk

We estimate the parameters of the risk process described in (1), namely, α_1, α_2, and β_1, by running patient-level regressions in two-steps. As a first step, we estimate an OLS regression of the next year's expenditures on the current information set:

$$c_{1,i} = \mu_1 + \gamma X_i^t + \varepsilon_{1,i} \tag{6}$$

where X_i^t includes patient characteristics potentially predictive of future costs, such as, cost, demographic information, and former utilization that are known initially. To study how predictability varies with time, we "normalize" $\alpha_1 = 1$ by defining the type θ_0 as the demeaned forecast:

$$\hat{\theta}_{0,i} = \hat{\gamma} X_i^t - \frac{1}{N} \sum_i \hat{\gamma} X_i^t \tag{7}$$

This is the empirical counterpart of the private risk type θ_0 introduced in the model above.[12] For robustness, we tried different definitions of X_i^t, ranging from cost-only to detailed utilization, demographics, and the presence of any of multiple chronic conditions. Appendix Table A1 summarizes the different sets of predictors we considered.

As a second step, we estimate α_2, capturing the decline in predictability due to mean-reversion, by regressing expenditures in two years' time on initial risk type $\hat{\theta}_{0,i}$ and the interim cost shock $\hat{\varepsilon}_{1,i}$:

$$c_{2,i} = \mu_2 + \alpha_2 \hat{\theta}_{0,i} + \beta_1 \hat{\varepsilon}_{1,i} + \varepsilon_{2,i} \tag{8}$$

Since we normalized $\hat{\alpha}_1 = 1$ we expect $\hat{\alpha}_2 < 1$, such that risk predictability falls with the horizon. We also expect $0 < \hat{\beta}_1 < 1$, such that interim cost shocks display some degree of persistence. Whenever a longer panel is available, the same method can be extended to estimate risk predictability at horizons longer than two years.

Note that this specification of risk does not explicitly account for moral hazard, i.e., that plan coverage directly affects spending. However, this effect leads, if anything, to *underestimate* mean-reversion. Indeed, moral hazard induces persistence (higher α_2) in this setting: after a high realization of first-year cost c_1, one can switch to a more generous plan. Moral hazard thus pushes cost c_2 in the second year upwards, reducing mean-reversion.

[12] For now, we assume risk is one dimensional, although this could be generalized (say, to distinguish between chronic and transitory conditions).

3.3 Results

Tables 1 presents estimates for the decline in risk predictability between one-year and two-year horizons. Across all specifications and for both samples, the estimated coefficient $\hat{\alpha}_2$ is close to 0.85, and statistically significantly different from 1. That is, at a two-year horizon, risk predictability is only about 85% of that at a one-year horizon. The two settings we study differ in the scope of potential moral hazard, so the fact that the coefficients are similar suggests that moral hazard has little influence on our estimates. The coefficient capturing the persistence of interim health shock, $\hat{\beta}_1$, is about 0.4 in the MarketScan data and 0.6 in the older Medicare sample.

Risk predictability declines over different horizons for the Medicare sample, where we have five data years, 2008–2012.[13] Table 2 shows estimates of the magnitude of the decline, normalizing one-year predictability to 1. Clearly, both the coefficients and the goodness of fit decrease over time. The constant decrease in predictability over time suggests the difference we focus on—between one-year and two-year horizons—further generalizes to longer periods.

4 Counterfactuals Under a Contract Length Reform

4.1 Approach

The previous section estimated how risk predictability varies with the horizon. In this section, we use the empirical distribution of risk from MarketScan claims to calibrate a model of equilibrium in insurance markets and derive counterfactual levels of coverage and welfare under two-year contracts. We estimate supply from risk data and derive demand assuming individuals face this same risk and maximize expected utility given their private information. We calibrate risk aversion to fit the observed fraction of individuals with High coverage in equilibrium with one-year contracts. We then contrast coverage and welfare with the alternative equilibrium, where contract horizon is extended to two years.

This approach is similar in spirit to Handel et al. (2015) among others. In principle, a more direct approach to estimate demand is to exploit quasi-random variation in contract terms and choice. However, the relevant variation for this paper does not exist in the U.S.: all contracts have the same horizon of one year. Instead, calibrating a model of insurance choice has the benefit of requiring only data on risk, as estimated in the previous section.

[13] For comparability across horizons we restrict the sample to only include individuals who were not covered throughout the period observed. (This restriction excludes some attrition due to mortality, but comparing the estimates with Column 1 of Table 1 shows the impact of the sample restriction on the estimate is negligible.)

Table 1: Predictability: Two-Year Horizon, Different Predictors

	Dependent Variable: *Log Future Medical Spending*		
	Information Set		
	I1	I2	I3
	A. MarketScan Sample		
α_2	0.842	0.868	0.871
	(0.001)	(0.001)	(0.001)
β_1	0.401	0.381	0.374
	(0.001)	(0.001)	(0.001)
R Sqr.	0.294	0.304	0.307
N	3,468,253	3,468,253	3,468,253
	B. Medicare Sample		
α_2	0.857	0.857	0.862
	(0.00140)	(0.00139)	(0.00135)
β_1	0.644	0.643	0.631
	(0.00193)	(0.00193)	(0.00195)
R Sqr.	0.621	0.621	0.623
N	456,482	456,482	456,478

Notes: Prediction of model (1) with log annuitized healthcare expenditure as the risk measure, and different predictor sets, defined in Table A1.

Table 2: Predictability Over Different Horizons

	Dependent Variable: *Log Future Medical Spending*		
	Prediction Horizon		
	2 years	3 years	4 years
α_2	0.858	0.772	0.686
	(0.00185)	(0.00209)	(0.00226)
β_1	0.640	0.543	0.472
	(0.00248)	(0.00266)	(0.00277)
R Sqr.	0.618	0.503	0.407
N	235,927	235,927	235,927

Notes: Prediction of model (1) with log annuitized healthcare expenditure as the risk measure, and different prediction horizons, for a balanced panel of Medicare beneficiaries over 2008–2012.

The downside is that is that one must make a parametric assumption about the utility function.[14]

Supply, discussed in Section 2, is estimated from the sample risk distribution by assuming premiums are fair and competitive, and therefore equal average cost. That is:

$$\Delta AC(p) = \iota_H \int_{\theta^*(p)}^{\infty} C(\theta)dF(\theta) - \iota_L \int_{-\infty}^{\theta^*(p)} C(\theta)dF(\theta) \tag{9}$$

The distribution $F(\theta)$ is estimated using its sample counterpart, the empirical CDF of the fitted values $\hat{\theta}$ from (7). The key observation is that the mapping from type to cost $C_t(\theta)$ varies with the horizon t, as specified in (1). For each value of θ, we estimate the distribution of $C_t(\theta)$, for $t = 1, 2$ using Bootstrap. Specifically, we draw with replacement 1,000 instances of ε_1 and ε_2 from the empirical distribution of the residuals from the estimation of (1a) and (1a) and use our estimates of $\hat{\alpha}$ and $\hat{\beta}$ to calculate the risk realization in each instance.

In the absence of exogenous variation in contract length (all contracts in our data are annual), we derive demand by assuming individuals choose the coverage level $b \in \{L, H\}$ that maximizes their expected utility: $E[P_b - OOP_b(C_\theta)|\theta] - \frac{\lambda}{2}Var[OOP_b(C_\theta)|\theta]$, where P and OOP are the plan premium and out-of-pocket spending for the cost $C(\theta)$. The mean-variance utility function specified here can be replaced by any utility that satisfies the single-crossing property. In our main specification, we assume that cost sharing is proportional to cost: $OOP_b(C_\theta) = \iota_b C_\theta$. Appendix Section A.2 shows the analysis for the case with out-of-pocket limits. Denote by $\theta^*(p)$ the type that is indifferent at this price difference between the plans. The aggregate demand is:

$$Q(p) = \int_{\theta^*(p)}^{\infty} dF(\theta). \tag{10}$$

We calibrate λ by fitting the equilibrium coverage, $Q(p)$ at which $\Delta AC(p) = p$, to the observed coverage with one-year contracts. We then compute counterfactual equilibrium where all plans are sold for two periods.[15] With two-period contracts, we discount both cost and utility using a common discount factor $\delta = 0.98$.

4.2 Results

Our counterfactual analysis suggests that increasing the contract horizon would increase equilibrium coverage by 6% and reduce deadweight loss by about 10%. The empirical distribution of $\hat{\theta}$ is used to estimate the supply and derive the demand. Figure 5 shows the

[14] See Aron-Dine et al. (2012) for a case in which agents' start of insurance coverage is staggered across the calendar year, introducing a *de facto* small variation in the contract horizon.

[15] In principle, the same exercise could be done for contracts longer than two years, at the cost of requiring more extrapolation.

empirical counterparts to the schematic figures from Section 2 obtained from one of our countrafactual simulations of the market. Panel (a) shows the market equilibrium with the baseline scenario of one-year contracts, with risk aversion calibrated to fit the national share of High coverage, 10%. Panel (b) shows the counterfactual equilibrium in this market when contracts are instead sold for two years. As the marginal cost of coverage exhibits mean-reversion, both supply (dashed) and demand (in gray) flatten to meet at a higher level of coverage. Panel (c) juxtaposes the two cases: one and two-year contracts, on top of one another, showing the empirical counterpart of Figure 3. Both supply and demand change between the baseline and counterfactual scenarios (Panels d–f).[16] The range of coverage we study covers the support of the distribution of actual coverage in the ACA exchanges, which is concentrated around 10% (see Figure A3).

Although demand by the sickest for High coverage is reduced due to the reversion of their expected costs towards the mean, the demand for High coverage of the *marginal* type increases, and therefore coverage overall increase. Intuitively, the important channel is that with longer horizon, healthier people are more likely to seek greater coverage. Figure 6 summarizes the increase in coverage for different initial levels of coverage. Although the system is non-linear, the increase in coverage under the counterfactual, longer horizon is roughly linear in the initial, shorter horizon. That is, the coverage level is predicted to be 6% higher under two-year contracts.Table 3 (Columns 1–3) summarizes the increase in coverage for different initial levels of High coverage. With out-of-pocket limits, the average increase in coverage is somewhat stronger and averages 10% (Figure A4).

In line with our discussion in Section 2, the overall increase in coverage is a combination of demand and supply responses. For low levels of initial coverage, the overall increase in coverage due to an extended contract horizon comes mainly from the supply side, whereas for higher levels of initial coverage, it comes mainly from the demand size (Table A4).

We predict that increased coverage due to increased contract horizon will be associated with an overall reduction in deadweight loss (Figure 7). This welfare improvement, albeit modest, exists for a wide range of initial coverage levels.[17] Increasing the horizon is not favorable to everyone: the healthiest types—those that rationally still opt for lower coverage in equilibrium even with a longer horizon—suffer welfare loss, as they are then exposed to greater risk. Their losses counteract some of the gains of the marginal types who obtain

[16]Note here that the willingness to buy High coverage for the most risky types is significantly larger than most estimates in the literature. Introducing liquidity constraints or out-of-pocket limits (see Appendix) would reduce it by orders of magnitude without altering the main results. Moreover, our welfare measure captures the deadweight loss at the other end of the risk distribution, largely sidestepping this issue.

[17]However, the points along the curve in Figure 7 represents different demands, so this is not strictly speaking a comparative statics exercise.

Figure 5: Counterfactual Market Equilibrium with Longer Contracts

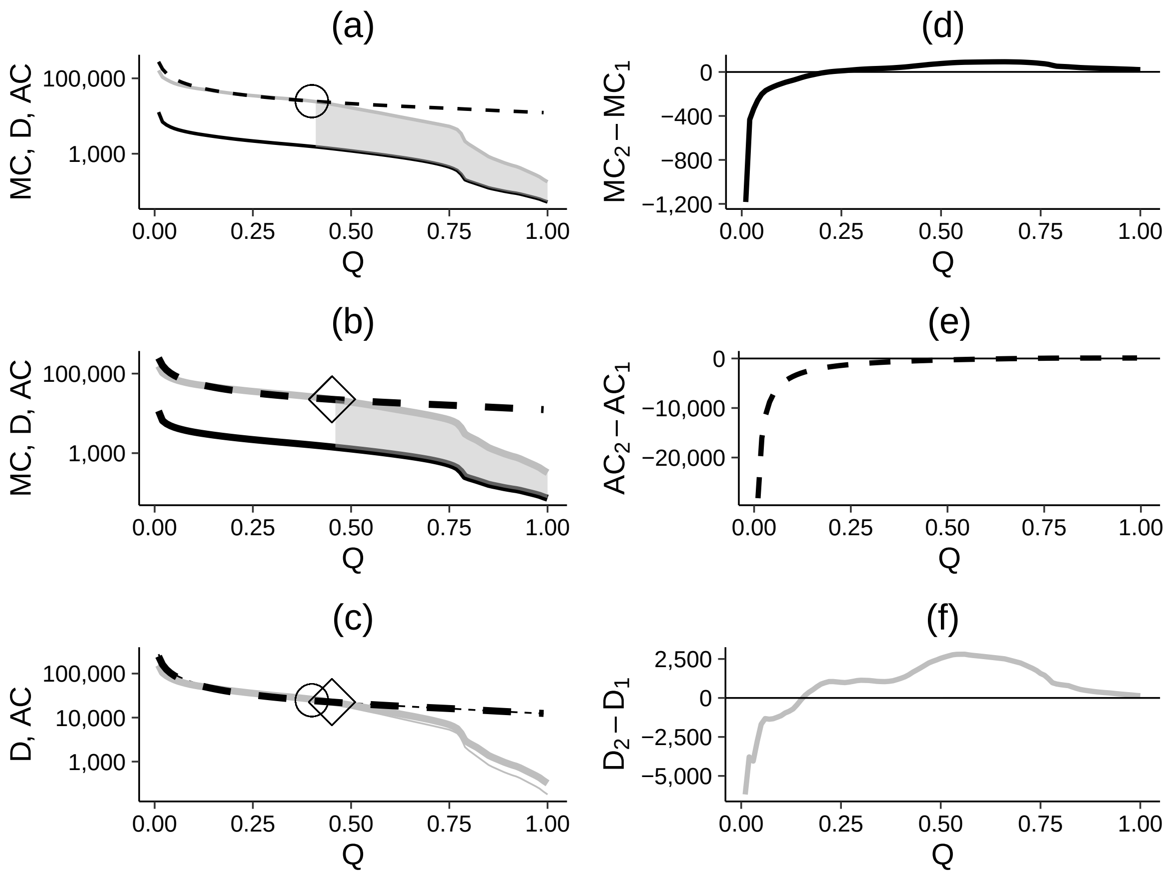

Notes: One year equilibrium (a) calibrated to match 50% High coverage and its two-year counterfactual counterpart (b). Demand (in gray) is derived from the marginal cost (solid black). The equilibrium coverage is at the intersection of AC (dashed) with D. Thicker lines depict counterfactual curves. Panel (c) combines both cases in one plot. Panels (d)-(f) show the change in each curve between equilibrium with one- and two-year contracts. We repeat this exercise for different levels of initial coverage.

Figure 6: Coverage Increase with Longer Contract Horizon

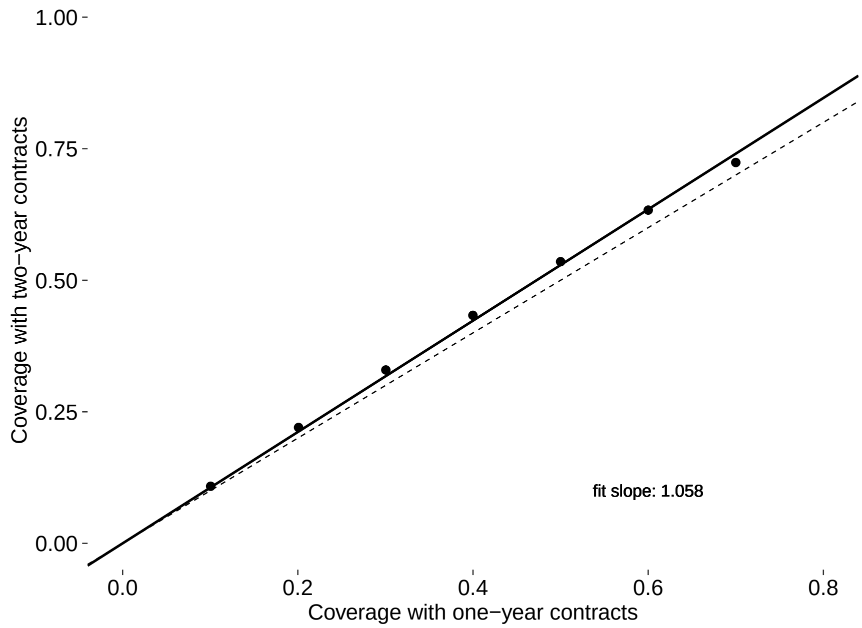

Notes: Share with High coverage in the baseline equilibrium with one-year contracts and in the counterfactual equilibrium with two-year contracts. The 45-degree line is dashed. Coverage increases roughly proportionally, to 6% above its initial level.

Figure 7: Deadweight Loss Reduction with Longer Contract Horizon

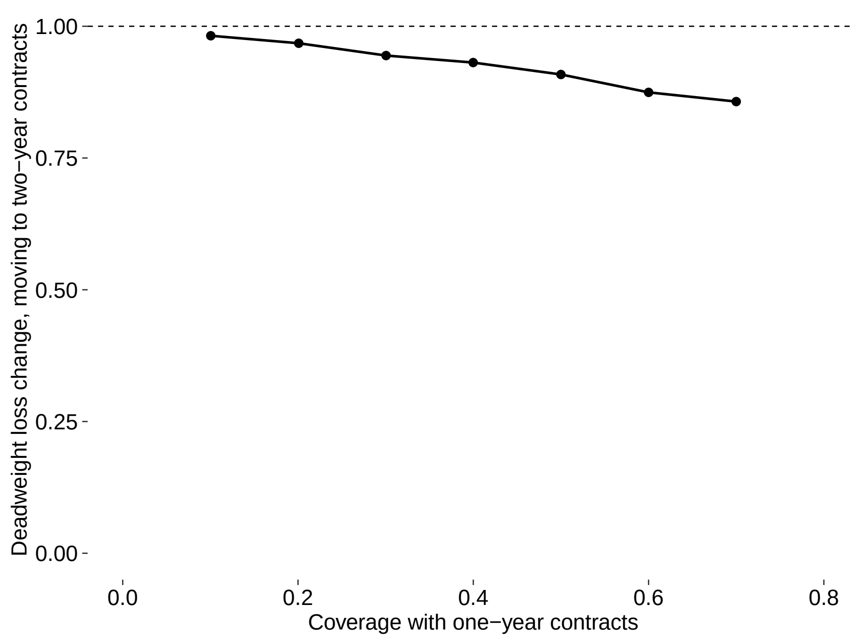

Notes: Welfare gains with longer contract horizon. The vertical axis shows the deadweight loss with two-year contracts as a fraction of its baseline level, with one-year contracts. Values below one (the dashed line) represent welfare gains. Gains increase with the baseline coverage.

Table 3: Horizon Effects on Coverage and Welfare

Baseline Coverage	Counterfactuals			
	Coverage Increase		Deadweight Loss Decrease	
%	p.p.	% relative to baseline	$	% relative to baseline
(1)	(2)	(3)	(4)	(5)
10	0.8	8.0%	−$97.00	−1.8%
20	1.9	9.7	−143.00	−3.2
30	2.9	9.7	−200.00	−5.6
40	3.4	8.4	−194.00	−6.9
50	3.5	7.1	−193.00	−9.1
60	3.3	5.6	−180.00	−12.5
70	2.4	3.4	−110.00	−14.3

Notes: Counterfactuals increase in coverage and decrease in welfare from moving to two-year contracts, for different levels of initial equilibrium coverage with one-year contracts. Coverage is the fraction of people buying High coverage. Deadweight Loss (DWL) is the average per-person forgone surplus (in dollars); it varies with baseline coverage. Negative values in columns (4) and (5) denote decreases in DWL.

High coverage only with a longer horizon, resulting in overall modest welfare gains. Overall, our results should be interpreted with this trade-off in mind, and while our risk estimates rely on a large sample of non-elderly Americans, it is plausible that different samples would yield negative results. However, our analysis provides tools to study it and can be replicated to obtain the magnitude of the welfare gains in other settings.

5 Discussion

The idea of extending contract length interacts with a number of issues related to health insurance markets.

First, two existing patterns of demand for insurance can influence the effect of contract length reform. Inertia in plan choice (as documented in Handel, 2013) implies that individuals might be slow to upgrade their coverage after moving to two-year contracts. However, one of the main results of the theoretical section is that the supply effect alone generates more coverage, even in the case in which the demand curve stays fixed. This result implies that inertia would reduce the efficiency gains of the reform, but not eliminate them.[18] Moreover, there is some evidence that consumers have different valuations for insurance contracts

[18] On the other hand, if inertia implies that agents do not react to plan prices, the reform will be ineffective, as lower costs do not translate to larger demand for greater coverage.

(Bundorf et al., 2012), i.e., that there is some amount of horizontal differentiation. These findings imply that contract length reform is particularly attractive as it preserves consumer choice. In particular, the components of insurance contracts that display the most preference heterogeneity (i.e., plan network) are largely orthogonal to horizon.

Another challenge to changing contract length in practice is commitment. In particular, moving to longer contracts involves restricting the possibility for individuals to change plans after a year, even for such individuals that incur greater out-of-pocket costs as a result. Currently, individuals can change their coverage around the year when any one of several qualifying life events occurs (such as marriage or change of employment status). Such exemptions could be abused to circumvent coverage change restrictions. While practical enforceability constraints are unlikely to be much more problematic for two-year contracts than they are with one-year contracts, they are important when considering contracts that extend further into the future. Handel et al. (2016) provides an in-depth discussion of long-term contracts with imperfect commitment.

6 Conclusion

This paper shows how extending the horizon of health insurance contracts impacts adverse selection in these markets. The main contribution of this paper is to show the implications of the dynamics of health risk *over time*, as opposed to simply its cross-sectional distribution. We show that increasing the contract horizon is another policy instrument that can be used to reduce selection. Conceptually, we argue that private information is endogenous to contract length because individual risk is harder to predict at longer horizons. This decrease in risk predictability is strongly borne by the data. Estimating a model of ACA-like exchanges, we find the effect of extending contracts to two years would be to expand coverage by about 6%. We also find positive, albeit moderate, welfare gains.

References

Aron-Dine, Aviva, Liran Einav, Amy Finkelstein, and Mark Cullen (2012), "Moral hazard in health insurance: How important is forward looking behavior?" NBER Working Paper.

Brown, Jason, Mark Duggan, Ilyana Kuziemko, and William Woolston (2014), "How does risk selection respond to risk adjustment? new evidence from the medicare advantage program." *American Economic Review*, 104, 3335–64.

Bundorf, M Kate, Jonathan Levin, Neale Mahoney, et al. (2012), "Pricing and welfare in health plan choice." *American Economic Review*, 102, 3214–48.

Cabral, Marika (2016), "Claim timing and ex post adverse selection." *The Review of Economic Studies*.

Cochrane, John H (1995), "Time-consistent health insurance." *Journal of Political Economy*, 445–473.

Culter, David M and Sarah J Reber (1998), "Paying for health insurance: The trade-off between competition and adverse selection." *Quarterly Journal of Economics*, 433–466.

Cutler, David M and Richard J Zeckhauser (1998), "Adverse selection in health insurance." In *Forum for Health Economics & Policy*, volume 1.

Diamond, Peter (1992), "Organizing the health insurance market." *Econometrica: Journal of the Econometric Society*, 1233–1254.

Einav, L and A Finkelstein (2010), "Selection in insurance markets: Theory and empirics in pictures." *The Journal of Economic Perspectives*, 25, 115–138.

Einav, Liran, Amy Finkelstein, Mark R Cullen, et al. (2010), "Estimating welfare in insurance markets using variation in prices." *The Quarterly Journal of Economics*, 125, 877–921.

Ellis, Randall P et al. (2000), "Risk adjustment in competitive health plan markets." *Handbook of health economics*, 1, 755–845.

Geruso, Michael (2013), "Selection in employer health plans: Homogeneous prices and heterogeneous preferences." Mimeo.

Hackmann, Martin B, Jonathan T Kolstad, and Amanda E Kowalski (2012), "Health reform, health insurance, and selection: Estimating selection into health insurance using the massachusetts health reform." *The American Economic Review*, 102, 498–501.

Hackmann, Martin B, Jonathan T Kolstad, and Amanda E Kowalski (2015), "Adverse selection and an individual mandate: When theory meets practice." *The American economic review*, 105, 1030.

Handel, Ben, Igal Hendel, and MD Whinston (2016), "The welfare impact of long-term contracts." Working Paper.

Handel, Ben, Igal Hendel, and Michael D Whinston (2015), "Equilibria in health exchanges: Adverse selection versus reclassification risk." *Econometrica*, 83, 1261–1313.

Handel, Benjamin R (2013), "Adverse selection and inertia in health insurance markets: When nudging hurts." *The American Economic Review*, 103, 2643–2682.

Heiss, Florian, Adam Leive, Daniel McFadden, and Joachim Winter (2013), "Plan selection in medicare part d: Evidence from administrative data." *Journal of Health Economics*, 32, 1325–1344.

Kolstad, Jonathan T and Amanda E Kowalski (2016), "Mandate-based health reform and the labor market: Evidence from the massachusetts reform." *Journal of health economics*, 47, 81–106.

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

Rice, Thomas, Marcia L Graham, and Peter D Fox (1997), "The impact of policy standardization on the medigap market." *Inquiry*, 106–116.

A Appendix

A.1 The CARA-Normal Example

A convenient tractable benchmark is to specify demand to be CARA over consumption, i.e., $u(C_t) = -exp(-\lambda C_t)$ for risk aversion λ, and health shocks $(\varepsilon_1, \varepsilon_2)$ to be iid $\mathcal{N}(0, \sigma)$. It highlights the same economics as our empirical analysis and has the advantage that it can be solved in closed form. In the CARA case, the net utility benefit of the High plan is given by:

$$v = -p + \Delta\iota(\mu + \tilde{\alpha}\theta_0) + \tilde{g} \tag{11}$$

where $\tilde{g} = \lambda/2((1-\iota_L)^2 - (1-\iota_H)^2)\tilde{\sigma}^2$ denotes the risk premium that individuals are willing to pay for insurance.[19] In this case, the key preference parameters are:

$$\tilde{\xi}_\theta = \tilde{\alpha}\Delta\iota \tag{12a}$$

$$\tilde{\xi}_0 = \Delta\iota\mu + \tilde{g} \tag{12b}$$

Demand rationally responds to horizon in two ways: (i) the sensitivity of v to private risk type $\tilde{\alpha}$ falls with the horizon as the predictability of future costs falls, and (ii) the risk premium $\tilde{g}$ increases with the horizon as consumers lose the option to adjust their plan choice after unexpected changes in their health status in interim periods.

In this case, the marginal type is given by:

$$\theta^*(p) = \frac{p - \Delta\iota\mu - \tilde{g}}{\Delta\iota\tilde{\alpha}} \tag{13}$$

[19] Risk is larger at a longer horizon because of persistence in interim shocks. At a one-year horizon $\tilde{\sigma} = \sigma$, whereas at a two-year horizon $\tilde{\sigma} = \sigma\frac{2+\beta_1}{2}$.

The equilibrium marginal type that determines equilibrium coverage solves $\theta^*(p) = AC(p)$ and has the following closed-form expression:

$$\theta^* = \theta_{max} - \frac{2(\tilde{g} - l)}{\tilde{\alpha}\Delta\iota} \tag{14}$$

This expression is intuitive. The share of the population choosing the High plan is proportional to $\theta_{max} - \theta^*$. This share increases with the gains from trade from the additional insurance of risk-averse agents $(\tilde{g}-l)/\Delta\iota$. Importantly for our mechanism, coverage decreases with the risk predictability $\tilde{\alpha}$ that governs the degree of selection on risk type in the market.

In the CARA-normal case, the welfare losses have an intuitive expression:

$$\begin{aligned} W^{loss} &= (\tilde{g} - l)\frac{\theta^* - \theta_{min}}{2\theta_{max}} \\ &= (\tilde{g} - l)(1 - S) \end{aligned} \tag{15}$$

where $\tilde{g} - l$ is the distance between the demand and MC curves, and $S = \frac{g-l}{\Delta\iota\alpha\theta_{max}}$ is the share of the population choosing the High plan.

Figure A1: Cost Response to Extended Contract Horizon

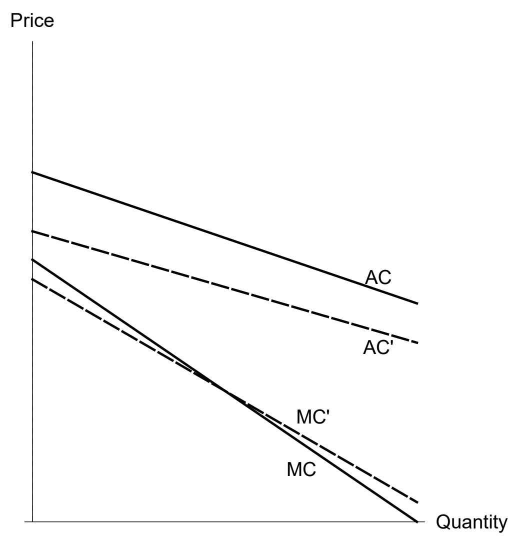

Notes: Average and marginal cost with one-year contracts (solid) and two-year contracts (dashed). Because of mean reversion, over longer horizons, between-individuals differences are smaller and more within-individual pooling is possible. Therefore, the difference in average cost between risk pools with High and Low coverage, AC, is smaller for all quantities.

Table A1: Risk Predictors Used in X_i^t in Different Specifications

Category	Predictors	I_1	I_2	I_3
Cost	total annual cost	v	v	v
Cost	cost by category (outpatient vs. inpatient)		v	v
Demographics	age, sex		v	v
Utilization	event counts, duration, and all other measures (see Table A3 for details)			v

Table A2: Choice of Coverage Level in the ACA Exchanges

Plan Level	**Actuarian Value**	**% Enrolled**		
Bronze	0.6	20%	*Low*	86%
Silver	0.7	65%		
Gold	0.8	9%	*High*	14%
Platinum	0.9	5%		

Notes: National average enrollment levels in the marketplace exchanges, by actuarial value. Actuarial Value is the average share of spending paid for by the plan. Source: CMS, April 2014. Excluding 1% catastrophic and unknown levels.

Figure A2: Empirical Cost Distribution

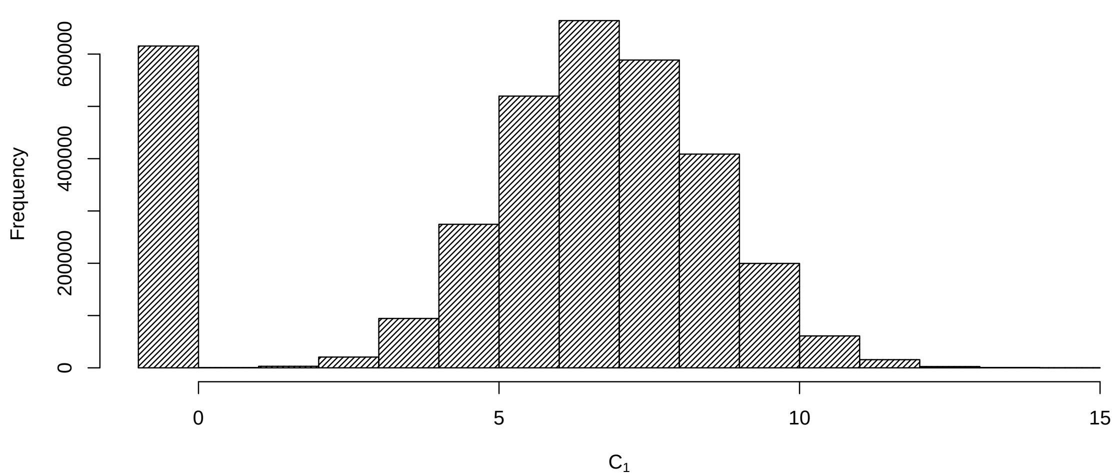

Notes: Histogram of the empirical distribution of log annual total healthcare costs (denoted c_1 in (1a)). The mass at zero had no expenditure during the year of coverage. Data source: MarketScan Research Databases.

Table A3: Descriptive Statistics: Cost and Use

	mean	sd
A. MarketScan		
Total Medical Costs	2,350	10,375
Total Inpatient Costs	757	7,474
Gender (Male)	0.465	0.498
Age	42.4	12.9
Inpatient Events	0.066	0.332
Outpatient Events	16.9	30.9
Observations	10,338,484	
B. Medicare		
Parts AB Total Costs	4,667	15,021
Part AB Annuatized Total Costs	6,003	23,968
Part D Total Prescription Costs	808	2,501
Gender (Male)	0.472	0.499
Age	66.29	1.710
Ambulatory Surgery Center Events	0.105	0.605
Part B Drug Events	1.395	5.516
Evaluation and Management Events	2.147	10.43
Anesthesia Events	0.181	0.714
Dialysis Events	0.0536	1.043
Other Procedures Events	2.387	9.841
Imaging Events	1.996	4.961
Tests Events	5.778	13.40
Durable Medical Equipment Events	1.050	4.485
Other Part B Carrier Events	0.690	6.705
Part B Physician Events	2.953	5.403
Part D Events	12.23	25.02
Acute Inpatient Covered Days	0.539	3.519
Other Inpatient Covered Days	0.115	2.065
Skilled Nursing Facility Covered Days	0.362	4.876
Hospice Covered Days	0.272	6.801
Hospital Outpatient Visits	2.658	11.20
Hospital Outpatient Emergency Room Visits	0.134	0.633
Inpatient Emergency Room Visits	0.0663	0.380
Home Health Visits	0.854	14.07
Parts AB Total Costs	4667.7	15021.1
Part AB Annuatized Total Costs	6003.3	23968.0
Part D Total Prescription Costs	808.1	2501.9
Log Annuatized Cost	5.294	3.792
Demeaned Log Annuatized Cost	-0.587	3.779
Part D Fill Count	16.02	30.16
Acute Inpatient Stays	0.118	0.527
Other Inpatient Stays	0.00798	0.115
Skilled Nursing Facility Stays	0.0150	0.179
Hospice Stays	0.00499	0.0805
Hospital Readmissions	0.0195	0.222
Observations	1,289,776	

Notes: Risk measures and predictors for both samples. See Section 3.1 for details.

Figure A3: Distribution of High Coverage

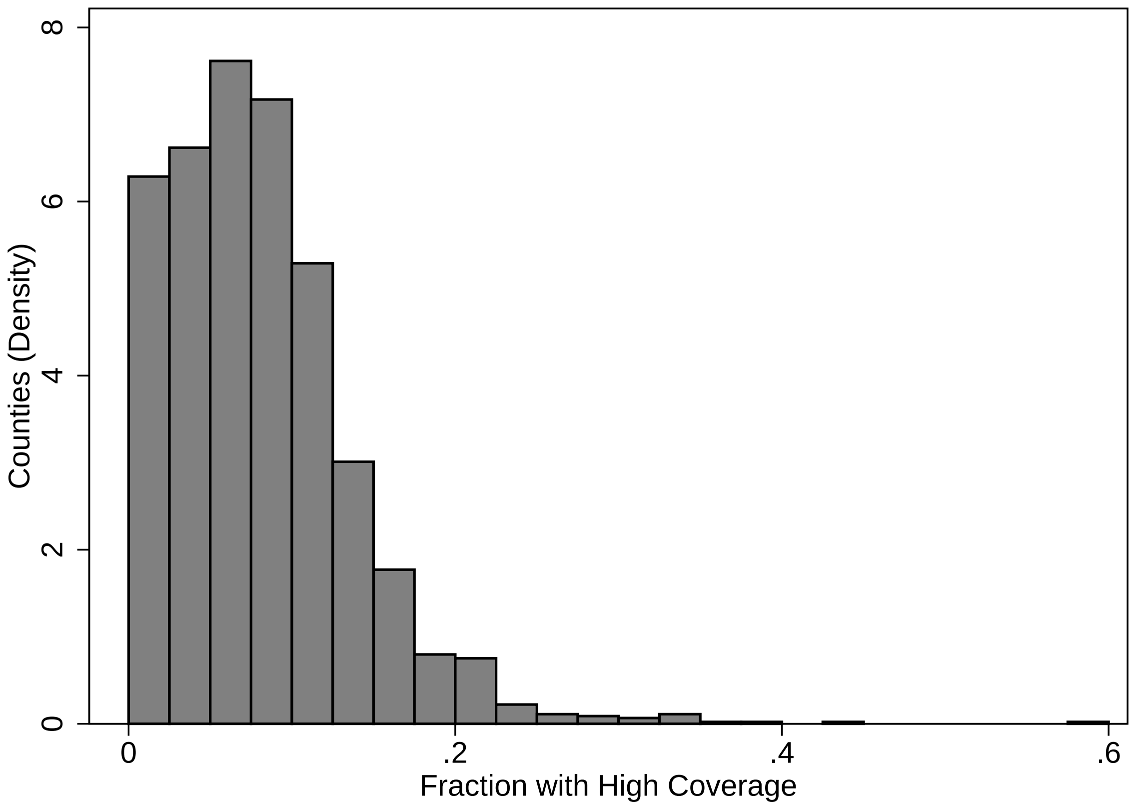

Notes: This histogram shows the distribution of county-level coverage for 1,807 counties in the 37 states with federal exchanges (counties are the smallest unit used for pricing at the ACA marketplaces). High coverage is defined as coverage with above Silver level (i.e., actuarial value of 80% or more). Source: Data on 2016 enrollment figures for qualifying health plan selections by metal level as of February 22, 2015, from data.cms.gov. Retrieved January 2017.

A.2 Counterfactuals with Out-of-Pocket Limits

This section repeats the counterfactual analysis of coverage and welfare with the additional assumption that contracts have an out-of-pocket spending limit. Out-of-pocket limits are currently mandated by the ACA.[20]

With an out-of-pocket limit M and a coinsurance rate of ι, the cost of coverage to an individual with cost C is $\min\{\iota C, M\}$; the rest of the cost is covered by the insurer. Introducing out-of-pocket limits that are uniform across plans violates the single crossing property. Namely, it is no longer true that types with higher expected risk benefit from High coverage more. For a small fraction of high-risk types whose total coinsurance payments are likely to be above the limit, and therefore similar between the Low and High plans, the Low plan is actually preferred to the High plan, as the premiums for the Low plan are lower. This violation is only true for a small fraction of the population, but complicates the analysis. We therefore abstract away from it, setting different out-of-pocket limits for the High and Low plans. Specifically, we pick $(M_H, M_L) = (1, 4) * 10^4$, which satisfies $\frac{M_H}{M_L} = \frac{1-\iota_H}{1-\iota_L}$. So, by construction, single crossing is preserved.

With out-of-pocket limits, moving to two-year contracts have similar estimated effects on coverage and welfare, although effects are somewhat larger in magnitude. Equilibrium coverage increases by 14% and deadweight loss decreases by 5–30%. Note, however, that when plans have out-of-pocket limits we only obtain results for initial coverage levels of .35 or higher, because in this case no level of risk aversion would generate lower initial coverage levels. Intuitively, out-of-pocket limits result in a flattened demand curve compared to the case of no such limits, because the utility from High coverage is not increasing as much in risk. Therefore, a larger fraction of high-risk types behaves similarly, either buying the High plan or not. Introducing further heterogeneity to demand could eliminate this problem.

[20] The mandatory out-of-pocket limit for a 2016 Marketplace plan is $6,850 for an individual plan and $13,700 for a family plan.

Figure A4: Coverage Increase with Longer Contract Horizon

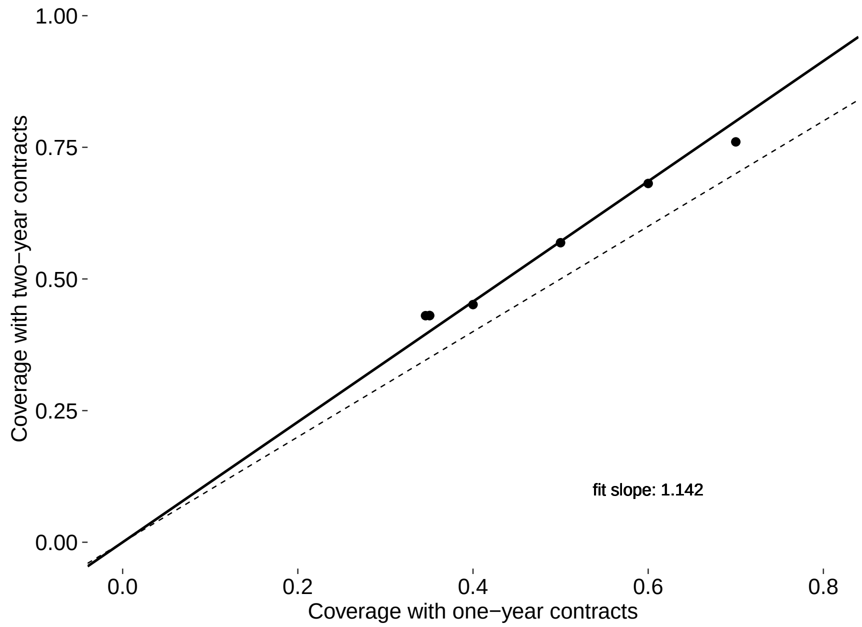

Notes: Share with High coverage in the baseline equilibrium with one-year contracts and in the counterfactual equilibrium with two-year contracts. The 45-degree line is dashed.

Figure A5: Deadweight Loss Reduction with Longer Contract Horizon

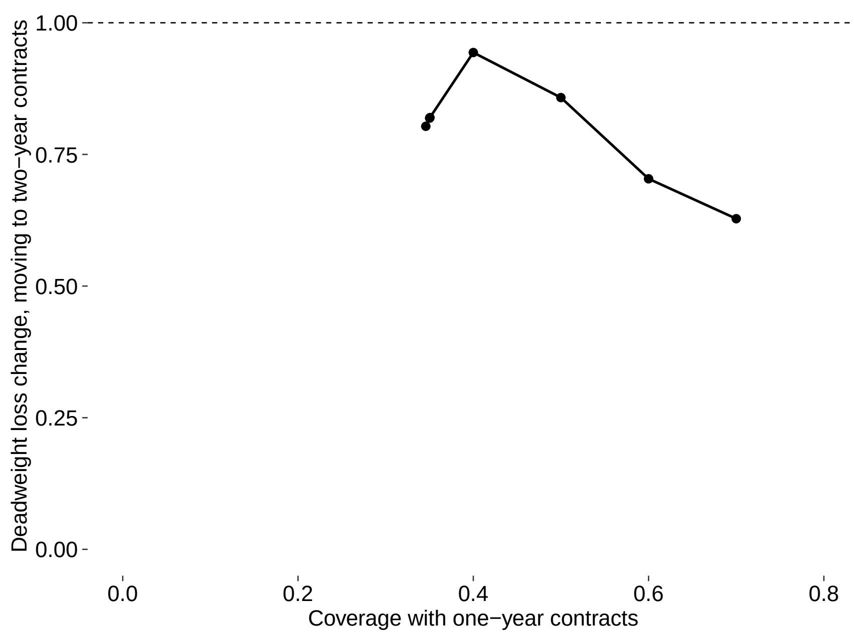

Notes: Welfare gains with longer contract horizon. The vertical axis shows the deadweight loss with two-year contracts as a fraction of its baseline level, with one-year contracts. Values below one (the dashed line) represent welfare gains. Gains increase with the baseline coverage.

Table A4: Supply-Side Contribution to the Overall Increase in Coverage

Baseline Coverage	Counterfactual Impact on Coverage: Supply Response		
	Change in Coverage: Net Supply Effects		
%	p.p.	% relative to baseline	% of overall response
(1)	(2)	(3)	(4)
10	0.9	8.9	112.5
20	0.8	4.2	42.1
30	0.7	2.3	24.1
40	0.4	1.1	11.8
50	0.2	0.5	5.7
60	0.1	0.2	3.0
70	0.0	0.0	0.0

Notes: Counterfactuals increase in coverage due only to supply-side changes from moving to two-year contracts, for different levels of initial equilibrium coverage with one-year contracts. Coverage is the fraction of people buying High coverage. For calculating the net supply-side effect of extended contract horizon on coverage, demand was left unchanged at its one-year (baseline) levels. Column (4) shows the the increase in coverage due to changes in supply relative to the overall increase due to changes in both supply and demand.