Horizon Effects and Adverse Selection in Health Insurance Markets

Olivier Darmouni and Dan Zeltzer*

May 15, 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, which leads 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 find that risk predictability falls significantly with horizon. Nesting these estimates in a simple equilibrium model of insurance markets, we find that a reform implementing two-year contracts would increase coverage by 6% from its initial level and yield average annual welfare gains of $100–$200 per person.

Keywords: Insurance markets, adverse selection, contract length

*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. 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.

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. 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 individuals with incomplete risk protection.[1]

This study examines how extending the length of health insurance contracts impacts adverse selection. The main contribution is to explore the benefits of increasing contract lengths as an instrument to limit market unraveling. Conceptually, longer contracts reduce selection because over longer horizons individual risk is less predictable. Lower predictability allows for better pooling of risk over time *within* individuals, as opposed to simply *across* individuals. 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. Nesting these estimates in a simple model of insurance exchanges suggests that increasing contract length from one to two years would increase coverage and reduce average deadweight loss in equilibrium.

We first show how to analyze an increase in 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 (e.g., by extending the time between open enrollment periods) 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 higher coverage if the mean type gets such coverage in the first place and towards lower 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.[2]

[1] In 2014, only 14% of enrollees chose one of the two high coverage plans, Gold or Platinum. By 2017, high coverage dropped to 5% (See Table A2). On a related note, premiums on these plans have been rising over the years.

[2] This paper focuses on the effects increasing contract length may have on adverse selection, without considering the issue of commitment. It is important to consider commitment for very long-term contracts, where it is hard to enforce (see Handel et al., 2016; Luz, 2015). However, we assume that enforcing commitment

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 populations. 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. Mean-reversion in risks suggests that longer contracts may limit the scope for adverse selection.

Finally, we provide some empirical evidence regarding the potential equilibrium effects 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 simple model of regulated insurance exchanges 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 by an average of $100–$200 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.

Practically, our results highlight how the annual periodicity of insurance enrollment in the United States has the implicit cost of increasing the scope for adverse selection. Some scope for selection is inherent in community-rated markets with free choice of coverage generosity, which include not only the ACA exchanges, but also the employer-sponsored health insurance market. We show that varying the frequency with which coverage can be adjusted is a new instrument for reducing such selection, by allowing for better risk-pooling within individuals. It adds to the existing and proposed policy instruments, which include enrollment mandates, risk adjustment, and premium penalty for lapsed coverage. These instruments all address a trade-off fundamental to the current healthcare policy debate in the United States—between choice and adverse-selection. We show that this trade-off also involves temporal aspect of consumer choice, and that the consequences of such choice are related to temporal properties

over only a slightly longer period than the current one is not substantially different.

of the underlying insured risk. Further understanding how consumers value the flexibility to adjust their coverage over time and their actual reaction to such flexibility are important questions for future research.

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; Ericson and Starc, 2016), risk adjustment mechanisms (Ellis et al., 2000; Glazer and McGuire, 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 why mitigating adverse selection through increased contract length could be helpful in preserving choice while mitigating selection.

However, only a few works 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), Cochrane (1995), and Luz (2015). This paper, in contrast, studies an environment similar to the current private U.S. insurance 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 less radical, and perhaps more realistic reform of existing marketplaces by looking at two-year, community-rated contracts.

The rest of this paper proceeds as follows. In Section 1 we study the impact of contract lengths on adverse selection in theory. In Section 2 we estimate a key determinant of this impact—the predictability of healthcare cost over different horizons. In Section 3 we nest these estimates in a parametric version of the model to study the counterfactual implications of longer contract length in equilibrium.

1 Mechanism: Contract Length and Adverse Selection

In this section, we present a simple framework that incorporates varying contract length into the standard model of insurance markets with adverse selection developed by Einav et al.

(2010) and used by Hackmann et al. (2012); Handel et al. (2015).

We define contract length as the time between two insurance plan choices. Concretely, a contract length reform can be implemented in two ways: either by regulating insurance products by extending their term, or by enforcing less frequent enrollment periods. For simplicity, we do not distinguish between these two implementations. Our focus is to consider contract length as a policy instrument and conduct a comparative statics exercise with respect to the regulatory environment. Concretely, we consider a reform whereby all contracts are sold for two years, instead of a one-year term.[3] We assume there is no co-existence of one-year and two-year contracts, nor are there hybrid contracts.[4]

We study selection among standardized plans of varying generosity that are community-rated, like in the ACA or the employer-sponsored insurance marketplaces.[5] The focus is on analyzing demand and supply response to contract length reform in order to study the effects of longer contracts on equilibrium coverage and welfare. The objective is to uncover a new potential benefit of extending contract length in community-rated markets. On the other hand, the potential issues associated with longer contracts, e.g. commitment, are well understood (see for instance Handel et al., 2016). Nevertheless, we make the plausible assumption that one-year and two-year contracts are equally enforceable.[6]

The mechanism we described is fundamentally different from the well-known idea that lifetime contracts chosen behind the veil of ignorance would remove adverse selection. Instead, we study a realistic potential reform that preserves three fundamental features of the current ACA health exchanges: consumer choice over plans, insurer competition, and community rating.

1.1 Health Insurance with Different Contract Lengths

Consider an insurance marketplace in which individuals periodically choose between two health insurance benefit plans—High and Low—with coverage levels $\iota_H > \iota_L$.[7] Denote

[3]The analysis extends to comparing short- and long-term contracts more generally.

[4]Sustaining an equilibrium with both long-term and short-term contracts seems difficult in practice. For a theoretical example using a dynamic pricing scheme, see Luz (2015).

[5]We choose to focus on the intensive margin of selection within coverage levels, conditional on buying insurance. But a similar logic applies to the extensive margin of selection into buying insurance.

[6]This is not to say that they are perfectly enforceable. Currently, individuals can change their coverage at any time during the year when any one of several qualifying life events occurs (such as marriage or change of employment). However, the incentives to abuse such exemptions are unlikely to be drastically different across one-year and two-year contracts. Handel et al. (2016) provides an in-depth discussion of long-term contracts with imperfect commitment. Given that enforcement frictions are particularly challenging to measure in our context, we leave the question of optimal contract length to future research.

[7]The analysis can be extended to more general contracts. We assume that regulated contract features such as co-insurance rates do not change with contract length, leaving the interaction between these policy instruments for future research.

by p the extra premium for purchasing High. Individuals maximize utility and have an underlying type $\theta_0 \in \mathbb{R}$, which determines their propensity to choose High over Low. This underlying type may depend on both individual risk (i.e., healthcare utilization costs) and other characteristics. When utility satisfies the single crossing condition, the aggregate demand for insurance is characterized by the marginal type $\theta^*(p)$ that is indifferent between the plans at price p, and is given by:

$$Q(p) = \int_{\theta^*(p)}^{\infty} dF(\theta). \tag{1}$$

We assume supply is regulated: premiums are community rated, rejections are prohibited, and coverage levels are standardized. The marginal cost of providing higher coverage to type θ^* is:

$$\Delta MC(p) = (\iota_h - \iota_l)C(\theta^*(p)), \tag{2}$$

where $C(\theta^*(p))$ is the expected cost of type θ^* during the contract term. To facilitate the comparison between one-year and two-year contracts, assume the premium difference p and costs are annualized. The difference in the average cost between the plans' risk pools is:

$$\Delta AC(p) = \iota_H \int_{\theta^*(p)}^{\infty} C(\theta)dF(\theta) - \iota_L \int_{-\infty}^{\theta^*(p)} C(\theta)dF(\theta). \tag{3}$$

In a competitive equilibrium, the premium difference p reflects this difference in average cost, i.e., $\Delta AC(p) = p$.

Adverse selection arises as costlier types are more inclined to purchase higher coverage. In such case, the marginal cost curve is downward sloping; the equilibrium price is generally higher than the efficient price, obtained when $\Delta MC(p) = p$. Namely, adverse selection leads to under provision of coverage.

Significantly, the key driver of the magnitude of adverse selection is the relevance of individual type for plan choice. Adverse selection is smaller when individuals with different risk types make similar plan choices.

This channel explains the effect of contract length reform on insurance demand and supply. In this setting, *risk predictability* will impact adverse selection. The less predictable risk is, the smaller the scope for selection, because the difference in expected risk $C(\theta_0)$ is smaller across different types θ_0. Graphically, when risk is less predictable, MC is flatter.[8] Moreover, risk predictability depends on contract length, because the time between plan choices determines the horizon over which predictions are to be made. Over longer horizons,

[8] In the extreme case of totally unpredictable risk, $C(\theta_0)$ is constant and therefore MC is fully horizontal. MC also equals AC in this case.

individual type θ_0 is a less effective predictor of future risk. As a consequence of this decline in risk predictability, both MC and AC curves are flatter with longer contracts (Figure 1).

Figure 1: The Effect of Increasing Contract Length on Cost

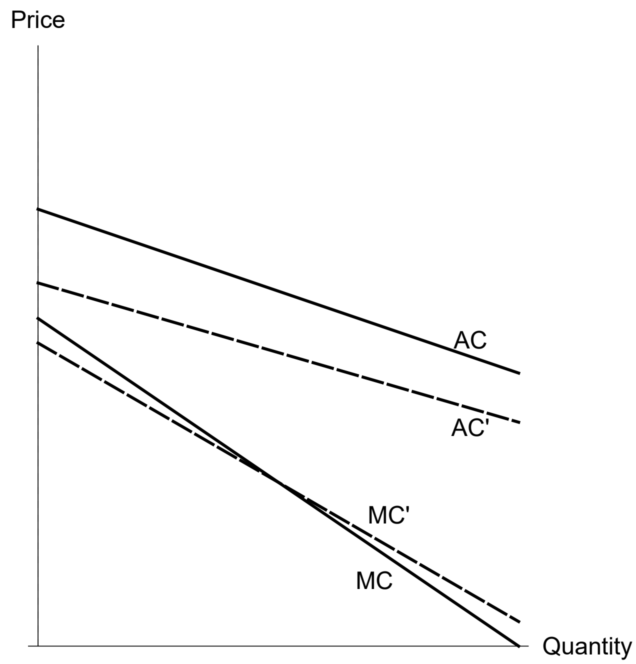

Notes: Average and marginal cost with one-year contracts (solid) and two-year contracts (dashed). Because of mean reversion, over longer horizons, between-individuals differences are smaller and more within-individual pooling is possible. Therefore, the difference in average cost between risk pools with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

This pivoting of the MC curve reflects an essential property of risk—mean reversion. Intuitively, over longer horizons, the expected risk of different types are slightly more like that of the mean type and therefore slightly closer to each other. We document mean-reversion in health expenditures in detail in the next section. In what follows, we highlight the effects of contract length on supply, demand, coverage, and welfare by focusing on the comparison between one- and two-year contracts. These horizons match our empirical estimates of risk predictability in the next section.

1.2 Equilibrium Effects on Coverage

We consider the effect of contract length reform, comparing two regulatory environments: one with one-year contracts, one with two-year contracts. We use our previous analysis to represent this comparative static exercise on one graph and argue that longer contracts change both consumer demand for insurance and its pricing by insurers.

The first main observation is that on the supply side, an increase in contract length increases coverage. Intuitively, coverage increases through a flattening of the AC curve.

Because AC shifts uniformly towards zero, this result is very general (Figure 2 illustrates this point). As contract length increases, mean reversion in risk leads to a flattening of the marginal cost curve and to a decline in the difference between risk pools for all coverage levels. 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. Under very general conditions, this change in the average cost curve induces an increase in equilibrium coverage.[9]

Figure 2: Supply Response with Extended Contract Length

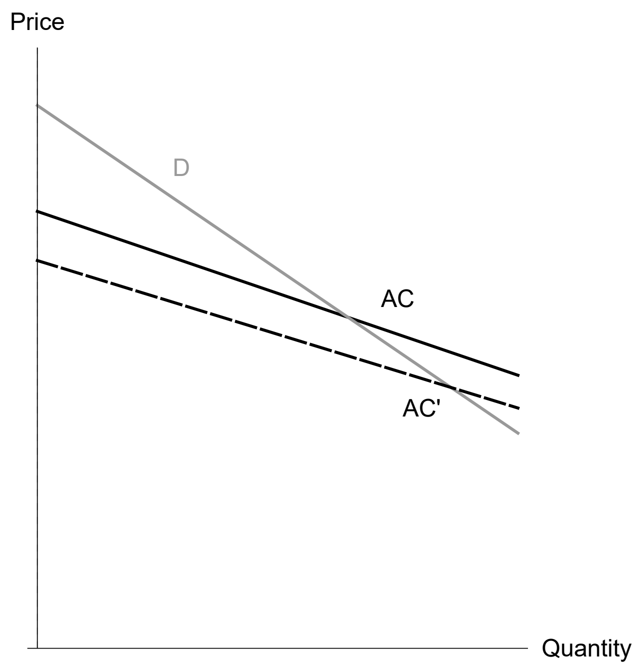

Increasing contract length reduces average cost by pooling risk over time, *within* an individual. This contrast with existing contracts that only pool risk *across* individuals. Longer contracts exploit the dynamics of risk: because risk displays some degree of mean-reversion, pooling risk over time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

Interestingly, the effect of contract length on demand can increase or decrease coverage. The intuition behind this result is illustrated in Figure 3. In both cases, an increase in contract length flattens the demand curve. 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 $\bar{\theta}$.

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

Figure 3: Demand Response and Equilibrium with Extended Contract Length

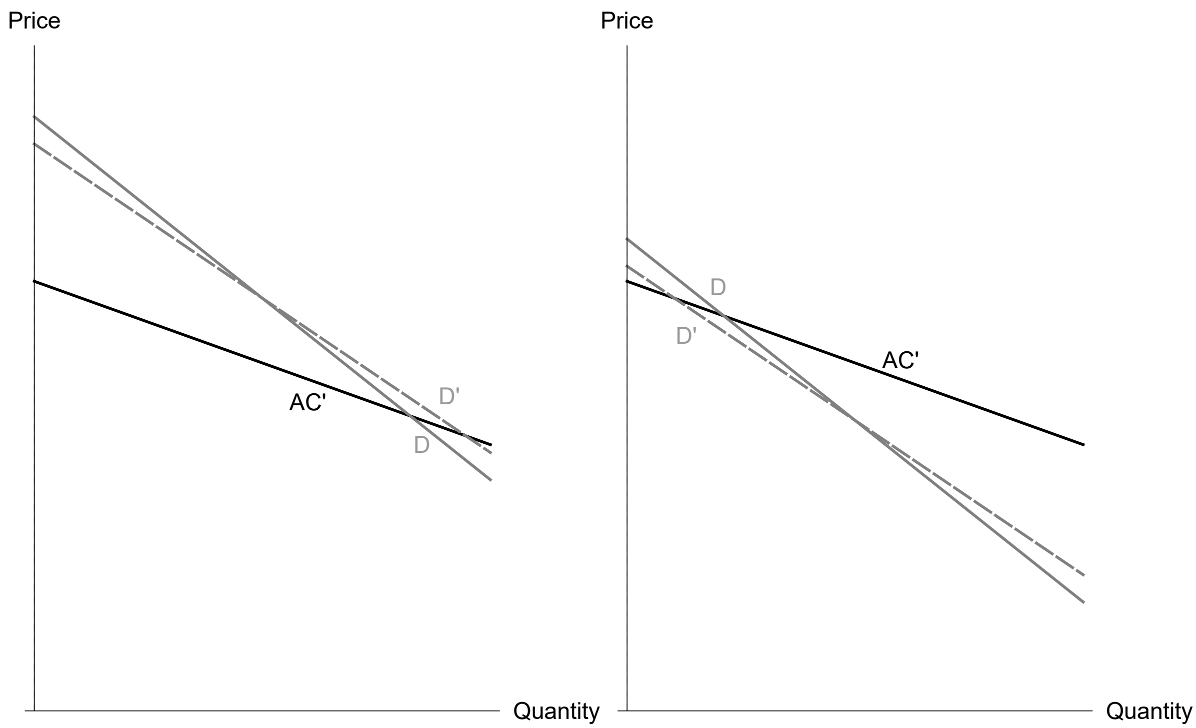

(a) $\theta^* > \bar{\theta}$ (b) $\theta^* < \bar{\theta}$

Notes: Equilibrium with one-year contracts (solid) and two-year contracts (dashed). 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).

Indeed, mean-reversion in health risk implies that types 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 a linear-uniform example, it would correspond to the average type $\mathbb{E}[\theta]$, but in general it can sit at any point of the risk distribution. The next section shows how to estimate it from the micro-data on health expenditures.

This explains the ambiguous demand channel. Effectively, as types become more homogeneous, individual plan choices tend to mirror the plan choice of the representative type. In the left panel of Figure 3, 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.

1.3 The Effects of Contract Length on Welfare

Figure 4 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. Indeed, they will face future interim shocks, which are unrelated to their current type θ_0, but (partly) persist until their next opportunity to upgrade to the High-coverage plan.

2 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.

Figure 4: Inefficiency Reduction due to Extended Contract Length

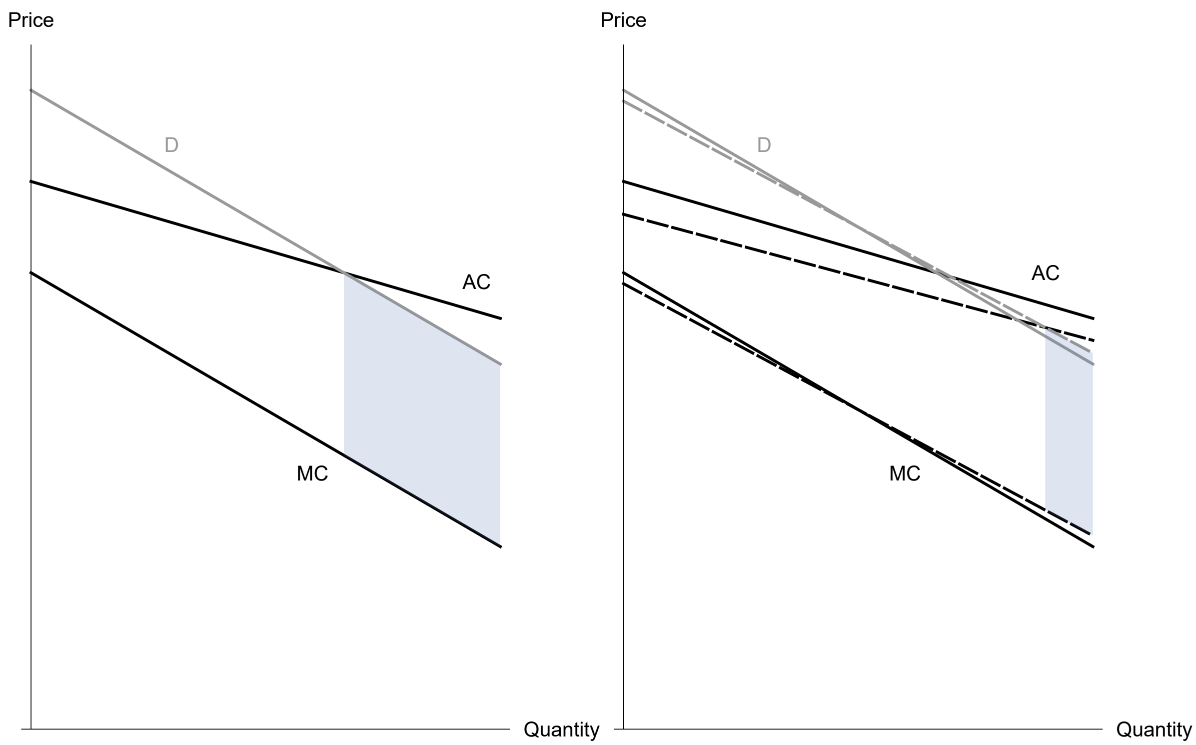

(a) One-year (b) Two-year

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 length leads to better within-individual pooling, which reduces selection and the deadweight loss associated with it.

2.1 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 of 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 5 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 of (4), presented below, show this reversion holds even when other predictors of risk beyond cost are included. Figure A1 shows the sample distribution of log cost c_1.

[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 5: Mean Reversion in Individual Healthcare Costs

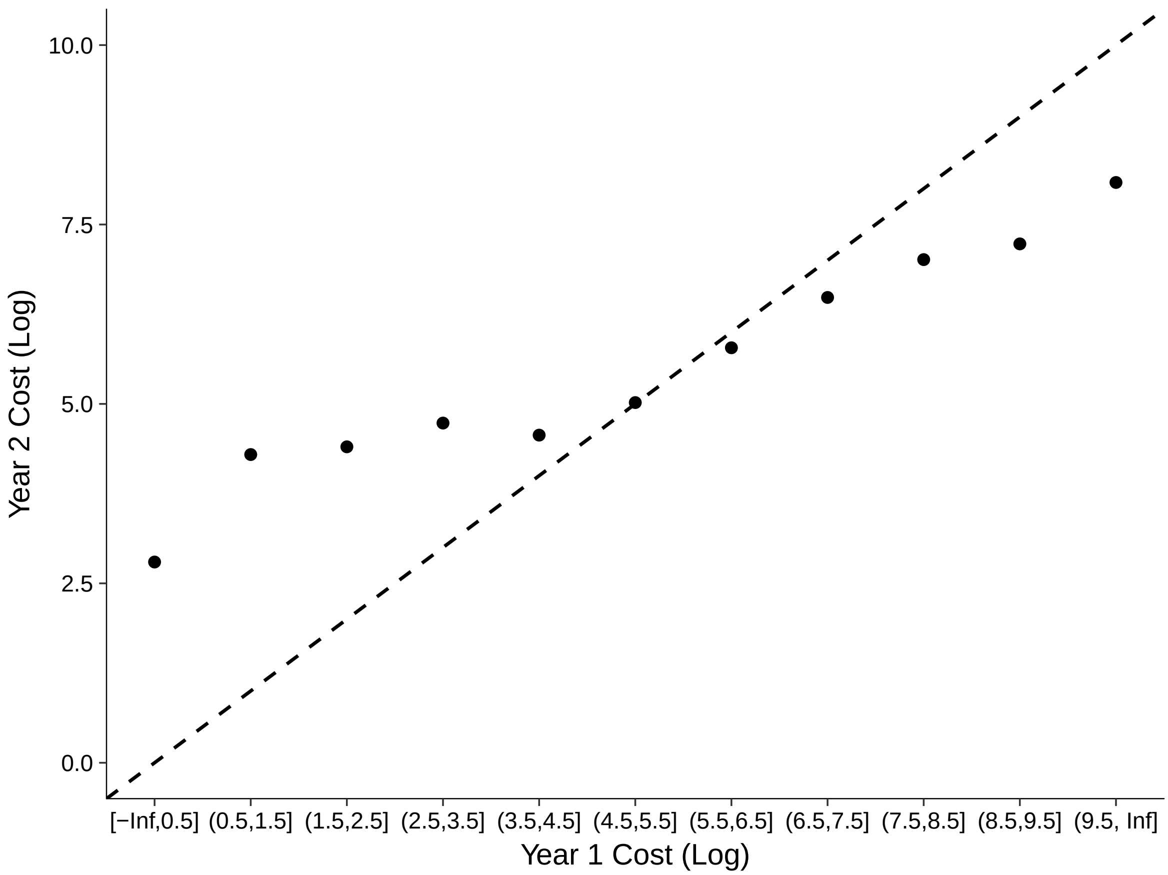

Notes: Binned scatter plot of individual healthcare costs over two subsequent years. Costs were aggregated at the individual level from administrative claims of 3.5 million working-age employees and their dependents from MarketScan Research Databases. The first-year costs were then binned into 11 equally 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 ($p<0.01$).

2.2 Empirical Model of Risk Predictability over Different Horizons

To compare the predictability of risk over different prediction horizons—one and two years—we study a finite risk process of the following log-linear form:[12]

$$c_1 = \mu + \alpha_1 \theta_0 + \varepsilon_1 \tag{4a}$$

$$c_2 = \mu + \alpha_2 \theta_0 + \beta_1 \varepsilon_1 + \varepsilon_2 \tag{4b}$$

where c is log individual cost (i.e., $C_{t,i} = e^{1+c_{t,i}}$ for $t = 0, 1, 2$), and θ_0, ε_1 and ε_2 are mutually independent, zero mean. We let individual type θ_0 be a function of both the baseline cost C_0 and other characteristics known at $t = 0$, which we denote by X_0. The key parameters determining the effects that changing the contract length has in equilibrium are α_1 and α_2. These parameters represent how predictive current risk type θ_0 is of future risks for next year and two years from now, respectively.

We estimate the model parameters using patient-level data in two steps. As a first step, we estimate an OLS regression of the next year's expenditures on the current information set:

$$c_{1,i} = \mu_1 + \gamma X_{0,i} + \varepsilon_{1,i} \tag{5}$$

where $X_{0,i}$ 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:

$$\hat{\theta}_{0,i} = \hat{\gamma} X_{0,i} - \frac{1}{N} \sum_i \hat{\gamma} X_{0,i} \tag{6}$$

This is the empirical counterpart of the private risk type θ_0 introduced in the general model above.[13] For robustness, we tried different definitions of $X_{0,i}$, 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 $\hat{\varepsilon}_{1,i}$:

$$c_{2,i} = \mu_2 + \alpha_2 \hat{\theta}_{0,i} + \beta_1 \hat{\varepsilon}_{1,i} + \varepsilon_{2,i} \tag{7}$$

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

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

Because 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 generosity directly affects spending. Moral hazard is less of a problem in Medicare, where coverage level is fixed. In the other context, this effect leads, if anything, to underestimating 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. The fact that the estimated decline in risk predictability turns out to be similar in these two settings that differ in the potential scope for moral hazard, suggests that moral hazard has in fact little influence on our estimates.

2.3 Results

Table 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 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, which is to be expected. It is surprising that while predictability increases when more comprehensive predictors are used, its decline over time remains similar across specifications and in different settings.

Risk predictability declines over different horizons for the Medicare sample, where we have five data years, 2008–2012.[14] 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.

[14]For comparability across horizons, we restrict the sample to include only 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 that 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		
	Total Cost	+Demog.	+Utilization
	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.

3 Counterfactuals Under a Contract Length Reform

The previous section estimated how risk predictability declines with the horizon, suggesting that the mechanism we highlighted in Section 1 is empirically plausible. Namely, increasing contract length, by extending the horizon over which predictions are made, should reduce the scope for adverse selection. In this section, we nest these estimates within a parametric version of the general model from Section 1 to provide preliminary evidence of what the magnitude of the effects of extending contract length on coverage and welfare might be in equilibrium.

Using the empirical distribution of risk from the claims data described in Section 2, we derive both supply and demand for insurance. We calibrate risk aversion to fit different initial levels of coverage in equilibrium in the baseline year. We then use this model to calculate counterfactual levels of coverage and welfare under an alternative regime with two-year insurance contracts. These estimates should be interpreted keeping in mind the following three caveats: (i) the risk data we use come from a specific population with employer-sponsored health insurance coverage; (ii) there is no exogenous variation in contract length to estimate demand response directly; (iii) our modeling choices are purposely stark. Nevertheless, we hope that these estimates provide some insight regarding the merits of extending health insurance contract length.

Setup We focus on comparing equilibrium outcomes between two cases: one- and two-year contracts. In both cases, individuals can periodically adjust their coverage level. In one case, they can do so every year.[15] In the other case, coverage can be adjusted only every other year. For comparability with one-year premiums, two-year premiums are annualized. 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. We set the standardized menu of two plans, High and Low, to be actuarially equivalent to the Bronze and Gold plans in the ACA exchanges, by setting the coinsurance rates to be 40% and 10%, respectively ($\iota_H = 0.9$ and $\iota_L = 0.6$).[16]

Supply To quantify the supply-side effects of moving to two-year contracts, we estimate the marginal and average cost curves (namely, ΔMC and ΔAC described in Equations 2 and

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

[16]We also analyze the case when plans also have out-of-pocket limits in the appendix .

3) for different contract lengths using the empirical risk process estimated from MarketScan Claims. The distribution $F(\theta_0)$ is estimated using its sample counterpart, the empirical CDF of the fitted values $\hat{\theta}_0$ from (6). We bin θ_0 into 100 bins (defined by its percentiles). For each value of θ_0, we estimate the distribution of $C_t(\theta_0)$, 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 (4a) and (4b) and use our estimates of $\hat{\alpha}$ and $\hat{\beta}$ to calculate the risk realization in each instance. Using this method, we calculate the cost curves for two cases: one-year contracts and two-year contracts, taking into account the difference in expected risks over different horizons. Figure 6 shows our baseline estimates of the supply side of the market for the one-year case, and Figure 7 shows the change in ΔMC and ΔAC due to moving to two-year contracts—the empirical counterpart of Figure 1.

Figure 6: Estimated Supply with One-Year Contracts

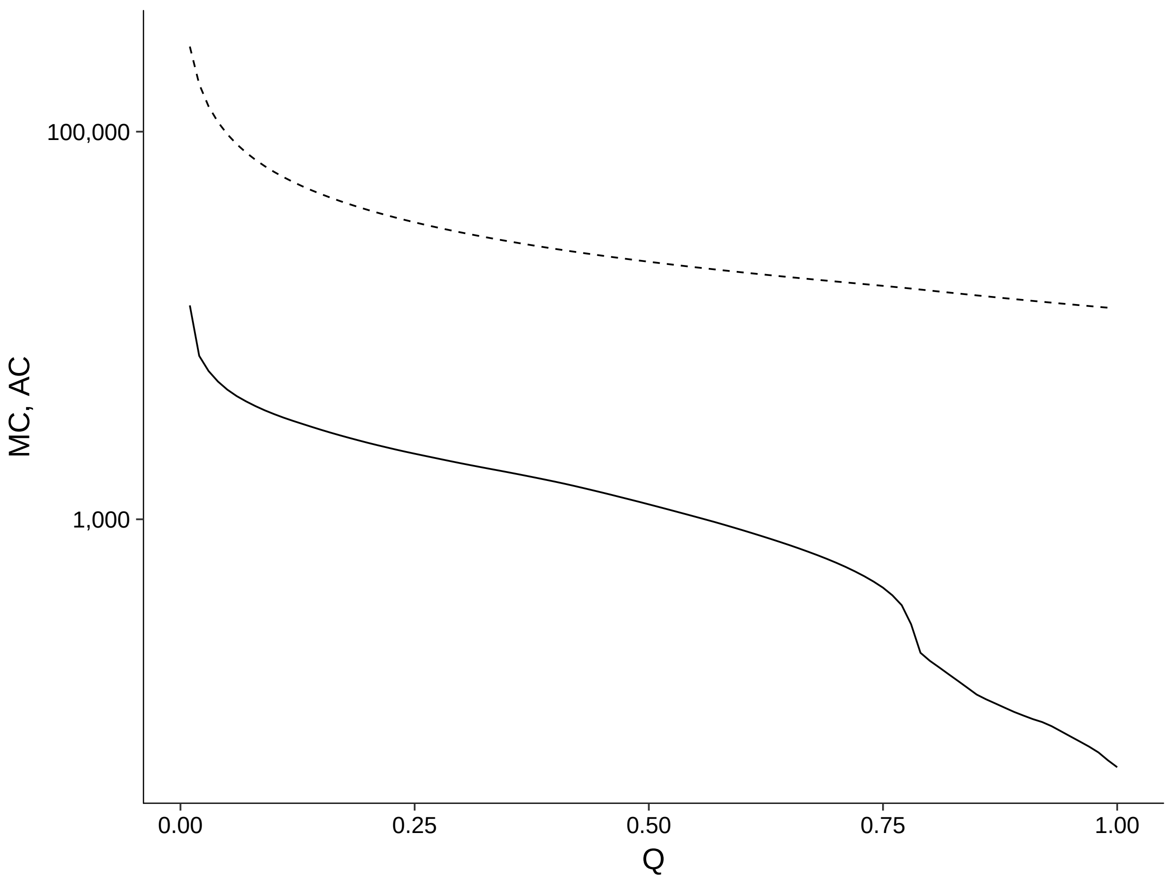

Notes: Estimated ΔMC (solid) and ΔAC (dashed), with one-year contracts. Curves are estimated using the empirical distribution of healthcare costs in MarketScan data. See main text for details.

Demand In the absence of exogenous variation in contract length (all contracts in our data are annual), we derive demand by assuming that individuals face this same risk process de-

Figure 7: The Changes in Cost and Supply with One-Year Contracts

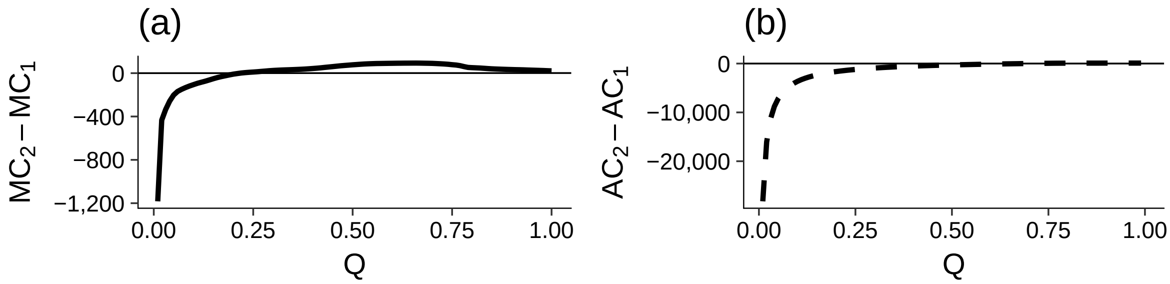

Notes: The change in ΔMC and ΔAC between one-year and two-year contracts, estimated using the empirical distribution of total healthcare cost in MarketScan data.

scribed in (4a) and (4b) and maximize expected utility. That is, we assume that individuals of type θ_0 choose coverage $b \in \{L, H\}$ to maximize expected utility $E[u(P_b - OOP_b(C_{\theta_0}))]$, where P and OOP are the plan premium and out-of-pocket spending for the cost $C(\theta_0)$. Our analysis uses a specific, mean-variance utility function: $E[P_b - OOP_b(C_{\theta_0})|\theta_0] - \frac{\lambda}{2} Var[OOP_b(C_{\theta_0})|\theta_0]$. The mean-variance utility function specified here can be replaced by any other utility function that satisfies the single-crossing property.[17] We discount both cost and utility using a common discount factor $\delta = 0.98$.

Equilibrium Our supply estimates are pinned down by claims data, but our demand estimates have one degree of freedom—risk aversion. We calibrate the risk aversion parameter λ to fit a range of initial levels of High coverage with one-year contracts. For each such initial level with one-year contracts, we compute the counterfactual equilibrium with two-year contracts, holding risk aversion constant. This yields a set of equilibrium pairs, allowing us to compare one- and two-year contracts for different levels of risk aversion (or, equivalently, for different levels of coverage with one-year contracts).

Welfare We measure the increase in coverage by comparing the one-year and two-year share of individuals who in equilibrium purchase (i.e., prefer, given equilibrium prices) the High plan over the Low plan. We measure the increase in welfare by calculating the decrease in deadweight loss. We define the deadweight loss, in the spirit of Einav et al. (2010), as the area between the demand curve and marginal cost curve. This area reflects unrealized beneficial trades by those buyers who value better insurance coverage more than its (private) marginal cost to them, but eventually give it up due to its high price.

[17] Implicitly, this specification of preferences assumes that individuals are rational and conduct the same prediction exercise as in the previous section. We leave extensions allowing for biases for future work, although interestingly, Section 1 shows that the supply response to contract length leads to an increase in coverage even if the demand curve does not change.

Remarks

In principle, it might be more direct to estimate demand ny exploiting 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. The downside is that one must make a parametric assumption about the utility function.[18]

One could use our same methods to study contracts longer than two years and even lifetime contracts (see Cochrane, 1995; Handel et al., 2016). We choose to focus instead on an incremental difference in contract length for two reasons. First, very long contracts present other complicating factors, such as aggregate uncertainty and commitment, that limit the feasibility of implementing them in practice. Second, from an empirical perspective, there is a greater confidence in extrapolating from the current annual contracts to an equilibrium with only slightly longer ones.

Results

Our counterfactual analysis suggests that increasing the contract length from one to two periods would increase equilibrium coverage by 6% and reduce deadweight loss by about 10%. Figure 8 shows the empirical counterparts to the schematic figures from Section 1 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 4. The changes between the baseline and counterfactual scenarios are driven by both a change in supply (Figure 7 above) and in demand (Panel d).[19] We repeat this exercise for different shifts in demand, corresponding to different levels of risk aversion and, respectively, to different levels of coverage by the High plan. The range of coverage levels we study covers the support of the distribution of actual coverage in the ACA exchanges across the United States (see

[18]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 length.

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

Figure A2).

Figure 8: Counterfactual Market Equilibrium with Longer Contracts

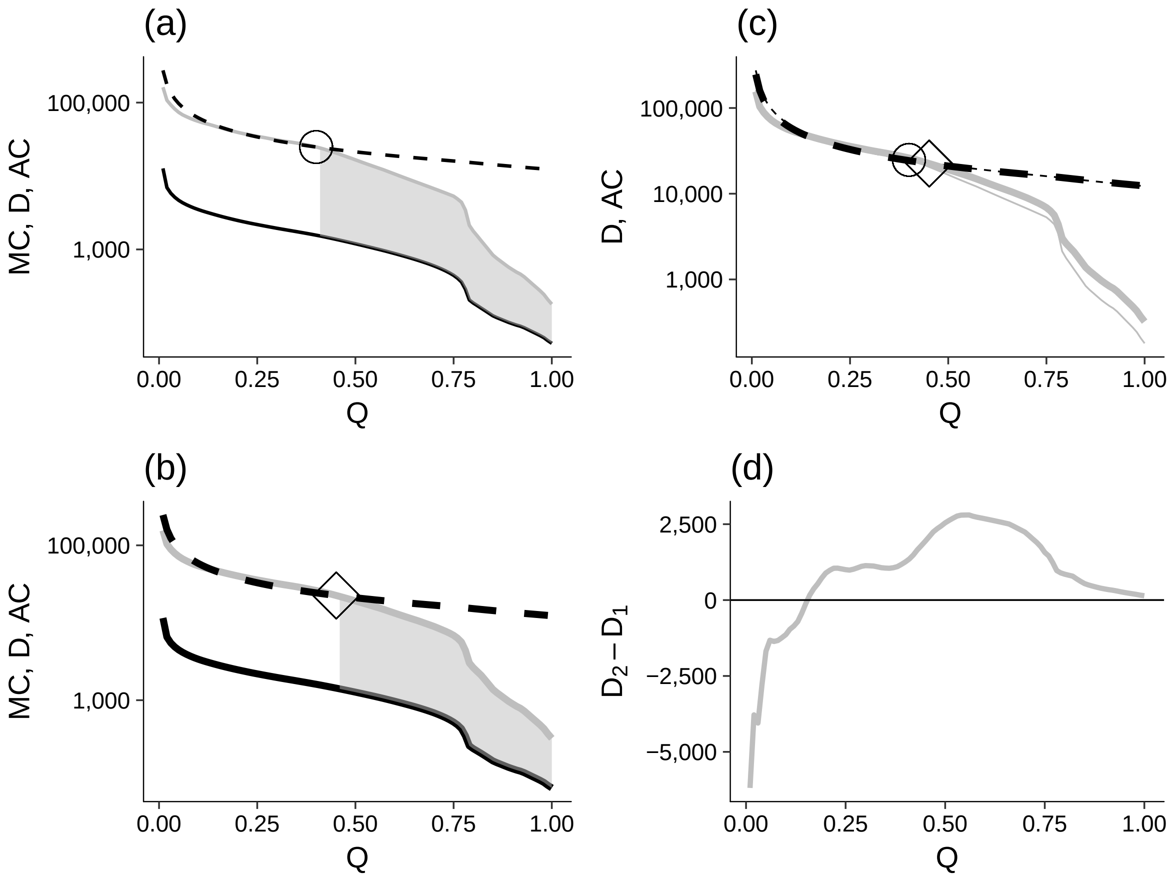

Notes: One year equilibrium (a) calibrated to match 50% High coverage and its two-year counterfactual counterpart (b). Demand (in gray) is derived from the marginal cost (solid black, plotted separately also in Figures 6 and 7 above). The equilibrium coverage is at the intersection of AC (dashed) with D. Thicker lines depict counterfactual curves. Panel (c) combines both cases in one plot. Panel (d) shows the change in demand between equilibrium with one- and two-year contracts. We repeat this exercise for different levels of initial coverage, which we match using different levels of risk aversion.

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 increases. Intuitively, the important channel is that with a longer horizon, healthier people are more likely to seek greater coverage. Figure 9 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 level of coverage. 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 A3).

Figure 9: Coverage Increase with Longer Contract Length

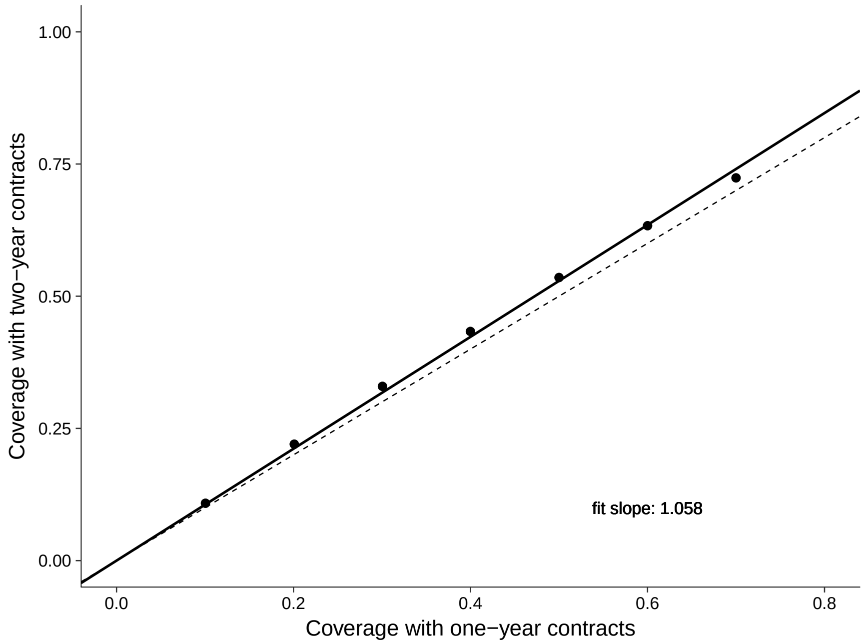

Notes: Share with High coverage in the baseline equilibrium with one-year contracts and in the counterfactual equilibrium with two-year contracts. The 45-degree line is dashed. Coverage increases roughly proportionally, to 6% above its initial level.

Figure 10: Deadweight Loss Reduction with Longer Contract Length

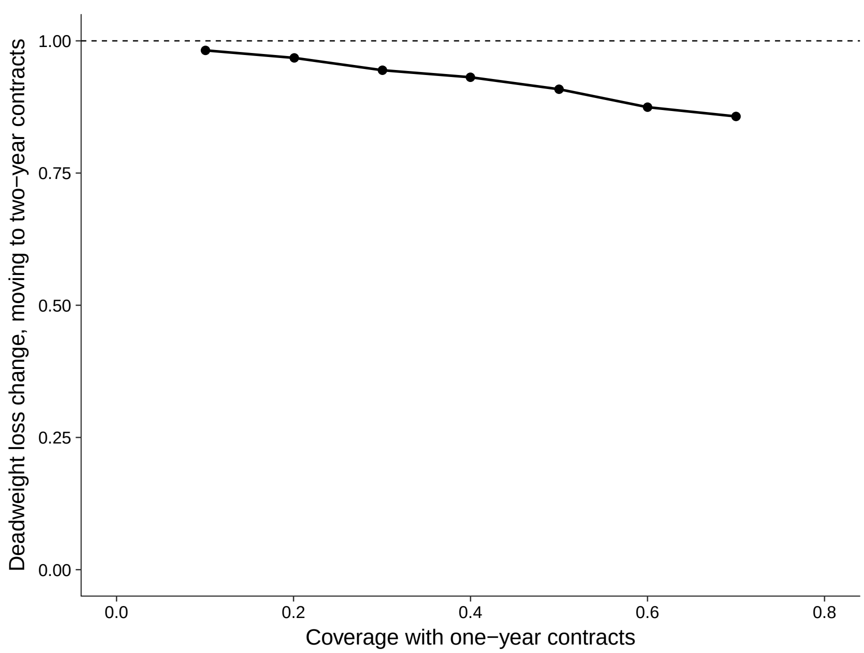

Notes: Welfare gains with longer contract length. The vertical axis shows the deadweight loss with two-year contracts as a fraction of its baseline level, with one-year contracts. Values below one (the dashed line) represent welfare gains. Gains increase with the baseline coverage.

In line with our discussion in Section 1, 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 length comes mainly from the supply side, whereas for higher levels of initial coverage, it comes mainly from the demand size (Table A4).

Table 3: Horizon Effects on Coverage and Welfare

	Counterfactuals			
Baseline Coverage	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.

We predict that increased coverage due to increased contract length will be associated with an overall reduction in deadweight loss (Figure 10). This welfare improvement, albeit modest, exists for a wide range of initial coverage levels.[20] 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 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.[21] 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.

[20]However, the points along the curve in Figure 10 represents different demands, so, strictly speaking, this is not a comparative statics exercise.

[21]In addition, there may be other implicit costs related to longer contracts.

4 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 length 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. Simulating 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, partly offset by the increased exposure of part of the population to greater risk.

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.

Ericson, Keith M Marzilli and Amanda Starc (2016), "How product standardization affects choice: Evidence from the massachusetts health insurance exchange." *Journal of Health Economics*, 50, 71–85.

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

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

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.

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.

Luz, Vitor Farinha (2015), "Dynamic competitive insurance." Discussion paper. University of British Columbia.

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.

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

Predictors	*Cost*	*Demographics*	*Utilization*
total annual cost	v	v	v
cost by category (outpatient vs. inpatient)		v	v
age, sex		v	v
event counts, duration, and all other measures (see Descriptive Statistics Table for details)			v

Table A2: Choice of Coverage Level in the ACA Exchanges

		Percent Enrolled			
Metalic Level	Avg. Coverage	2014	2015	2016	2017
Bronze or Silver	60%–70%	86%	90%	92%	95%
Gold or Platinum	80%–90%	14%	10%	8%	5%
Enrollment (Millions)		6.3	10.2	11.1	12.2

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, KFF. Excluding 1% catastrophic and unknown levels.

Figure A1: Empirical Cost Distribution

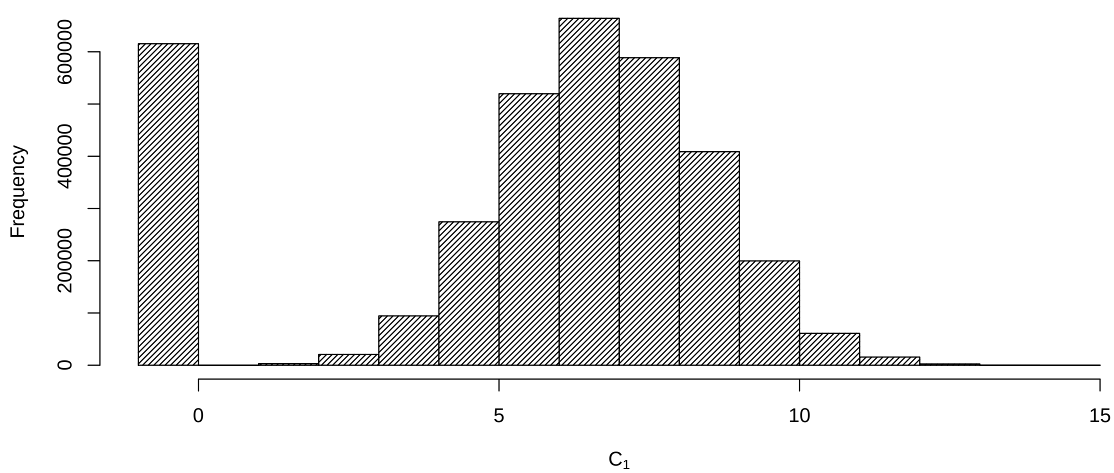

Notes: Histogram of the empirical distribution of log annual total healthcare costs (denoted c_1 in (4a)). The mass at zero had no expenditure during the year of coverage. Data source: MarketScan Research Databases.

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 2.1 for details.

Figure A2: Distribution of High Coverage

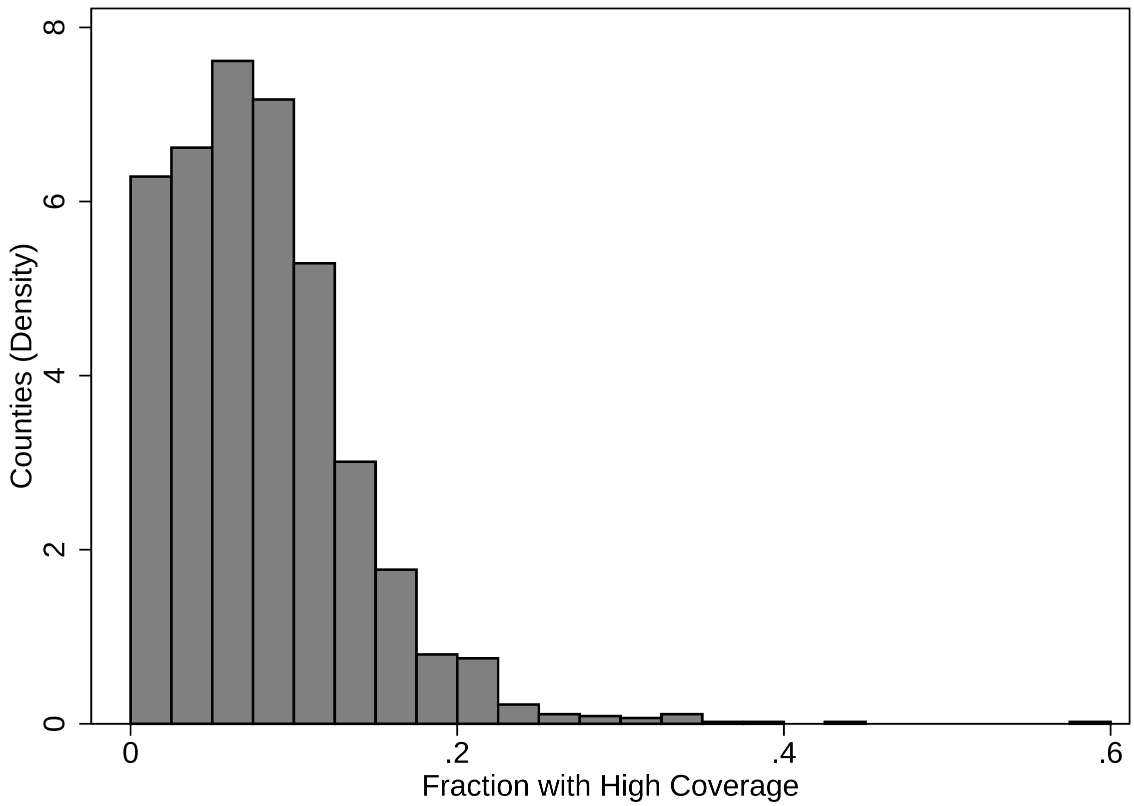

Notes: This histogram shows the distribution of county-level coverage for 1,807 counties in the 37 states with federal exchanges (counties are the smallest unit used for pricing at the ACA marketplaces). High coverage is defined as coverage with above the Silver level (i.e., actuarial value of 80% or 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 Appendix

A.1 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.[22]

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 more from High coverage. 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 has 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.

[22]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 A3: Coverage Increase with Longer Contract Horizon

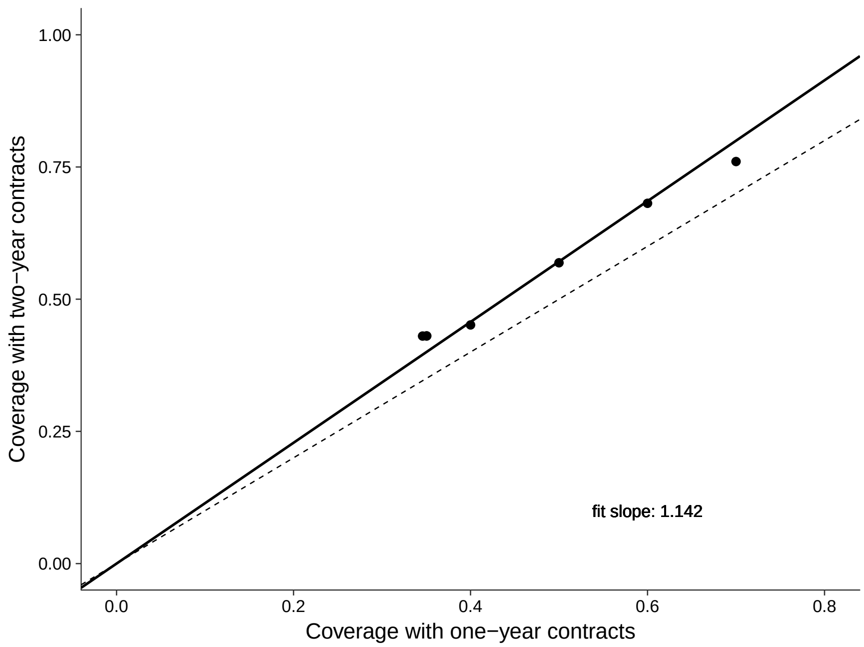

Notes: Share with High coverage in the baseline equilibrium with one-year contracts and in the counterfactual equilibrium with two-year contracts. The 45-degree line is dashed.

Figure A4: Deadweight Loss Reduction with Longer Contract Horizon

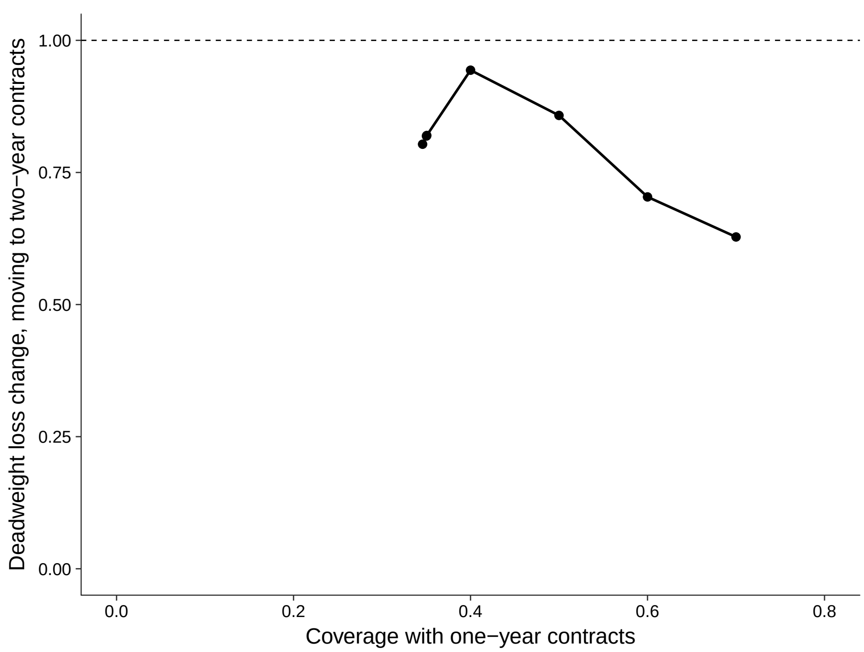

Notes: Welfare gains with longer contract horizon. The vertical axis shows the deadweight loss with two-year contracts as a fraction of its baseline level, with one-year contracts. Values below one (the dashed line) represent welfare gains. Gains increase with the baseline coverage.

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 increase in coverage due to changes in supply relative to the overall increase due to changes in both supply and demand.