

Article

The Shifting Shape of Risk: Endogenous Market Failure for Insurance

Thomas G. Koch

Bureau of Economics, Federal Trade Commission, Washington, DC 20580, USA; tkoch@ftc.gov;
Tel.: +1-512-809-8014

Academic Editor: Luca Regis

Received: 6 July 2016; Accepted: 9 December 2016; Published: 27 January 2017

Abstract: This article considers an economy where risk is insurable, but selection determines the pool of individuals who take it up. First, we demonstrate that the comparative statics of these economies do not necessarily depend on its marginal selection (adverse versus favorable), but rather other characteristics. We then use repeated cross-sections of medical expenditures in the U.S. to understand the role of changes in the medical risk distribution on the fraction of Americans without medical insurance. We find that both the level and the shape of the distribution of risk are important in determining the equilibrium quantity of insurance. Symmetric changes in risk (e.g., shifts in the price of medical care) better explain the shifting insurance rate over time. Asymmetric changes (e.g., associated with a shifting age distribution) are not as important.

Keywords: risk sharing; medical insurance; adverse selection

JEL Classification Codes: D3; I1

1. Introduction

This article considers how the price for insurance and the number of people insured varies when the risk and other characteristics of the population change. We study these relationships, known as comparative statics, for insurance markets that suffer from either adverse or advantageous selection, and then consider the specific case of medical insurance markets.

The comparative statics of medical insurance markets is of interest, as these markets leave millions of Americans uninsured against medical risk. Eighty percent of the private medical insurance in the US is employer-provided¹, and employers are not allowed to discriminate according to relevant demographic and medical information when they price insurance provided to their employees.² Policies, such as the income-tax deductibility of health insurance, or a “Cadillac tax” on expensive health plans, should induce changes in prices and quantity. I demonstrate that increases in the cost of insurance, such as the “Cadillac tax”, may actually increase the amount of insurance, rather than decrease it. The actual price and quantity effects of policies designed to shift either the supply or demand curves in insurance markets may have counter-intuitive effects. For example, many aspects of the Affordable Care Act (ACA), such as the Cadillac tax or ACA plan subsidies, that are designed to decrease or increase the amount of insurance, respectively, may have opposite effects.

In these models, the amount of risk left uninsured due to selection depends upon the distribution of risks—for example, what is the ratio of high-cost to low-cost individuals, and how different in average costs are high- and low-cost individuals. We use repeated cross-sections of medical expenditure

¹ This number is based upon the author’s calculations using the MEPS, which is described below.

² This is why a menu of contracts is unlikely to solve the adverse selection problem here: any separate pooling they could induce would rely on indirect, and possibly manipulable, correlates with expected cost.

data to infer such ratios, and use them in a model of insurance markets with asymmetric information. I also find that changes to the distribution of medical risk faced by consumers in recent years should have led to more expensive medical insurance. It should also have led to more people with medical insurance.

The price of employer-provided medical insurance to cover a single employee has grown five percent a year from 1996 to 2005. The price of family coverage has grown 5.7 percent per year over the same period.³ At the same time, more Americans have gone without medical insurance. Among the non-poor and non-elderly, the percentage of individuals without private insurance rose 22 percent (18 percent to 22) from 1996 to 2004.⁴

One explanation for higher insurance prices is the well-documented growth in average medical spending (Swartz [1]; Freudenheim [2]). As medical costs grow, the supply curve for medical insurance shifts up, and induces movement along the demand curve. This new equilibrium price is higher, while equilibrium quantity is lower. However, changes in the distribution of medical expenditures (e.g., mean, median, variance, etc.) are the result of changes in the distribution of medical risk. In this case, the demand curve also shifts up. The new equilibrium price is larger, but whether the new equilibrium quantity is larger or smaller is unclear.

This article attempts to disentangle these changes. This study complements the findings of Gruber and Levy [3], which documented similar (though differently measured) changes to the distribution of medical risk. We use changes in the distribution of medical expenditures in the US to infer changes in the distribution of medical risk. We then take these measured differences in the distribution of medical risk into a quantitative model of insurance choice with asymmetric information. We find that measured changes in the distribution of medical risk should have made insurance more expensive and more prevalent. While the former trend is consistent with the data, the latter is not.

This work can help shed light on recent research that considers alternative risk-adjustment factors; that is, what kinds of information could be used to price insurance that would limit the amount of market failure due to adverse selection. Both [4,5] find evidence that age-based health insurance prices may drive some of those age-specific markets to collapse. Here, we identify some of the characteristics of risk pools that make them sustainable.

The lessons from this exercise go beyond the sectoral concerns of medical insurance. If the source of uninsurance in other markets is asymmetric information, then this reminds us that distributions matter. If there is an across-the-board increase in risk, then insurance markets likely would provide more insurance, instead of less, as the demand curve shifts out further than the supply curve because of risk aversion. Furthermore, it introduces the notion that comparative statics do not depend upon the nature of the selection (adverse vs. advantageous). Instead, the comparative statics depend upon the stability of the equilibrium. For example, if wage volatility goes up, it matters whose wage volatility goes up.

This also relates to the market failure as discussed in recent work on insurance against aggregate, catastrophic risk [6–8]. When a natural disaster strikes, it may test the solvency of the risk pool, leading to partial insurance. The risks under study here have idiosyncratic realizations, as opposed to the widespread loss of property due to an earthquake or hurricane. For example, we do not consider how insurance markets might deal with a flu pandemic. Likewise, the tax subsidy of health insurance should mitigate adverse selection, and the indirect nature of the contract (typically via employers, instead of directly to consumers) may also influence the extent of adverse selection.

³ The level changes from 1996 to 2005 for these two prices is 1992 to 3227, and 4954 to 8675, as reported by the MEPS-IC. These calculations adjust for changes in the price level, as measured by the Urban CPI.

⁴ These numbers are calculated from the MEPS-HC, which collects information from a nationally-representative sample of households.

2. A General Framework with Initial Results

Suppose there is a unit measure of risk-averse agents, heterogenous in their financial risk, \tilde{x} . Each agent's risk is characterized by $\lambda \in \mathbb{R}$, with conditional expectation $\mathbb{E}[\tilde{x}|\lambda] \in \mathbb{R}$. Risk types are distributed according to a second distribution, $\lambda \sim \Gamma$. Agents with wealth w and preferences $u(\cdot)$ ⁵ have a willingness to pay for insurance, $\pi(\lambda)$, which is the dollar amount that makes the agent indifferent between paying $\pi(\lambda)$ for insurance, or facing the risk without insurance:

$$u(w - \pi(\lambda)) = \mathbb{E}[u(w - \tilde{x})|\lambda].$$

Risk aversion means that willingness to pay for insurance is greater than expected outcome. An agent purchases insurance if his or her willingness to pay for insurance is larger than its price, $\omega \in \mathbb{R}$ —that is, if $\pi \geq \omega$, and not otherwise. An inverse demand function, $\pi^{-1}(\omega)$, provides the risk type(s) associated with a particular level of demand at price ω .

The demand curve for insurance is simply the ordering of the individual demands:

$$D(\omega) = 1 - \Gamma(\pi^{-1}(\omega)).$$

The demand curve is downward sloping by construction.

We assume that insurance firms provide insurance via price competition. These firms cannot provide risk-type contingent contracts because of asymmetric information. The sources of this asymmetric information may either be based in the microfoundations of the model, or due to legal constraints imposed upon insurers. In the case of the US medical insurance market, the latter case is certainly true—firms that offer their workers medical insurance are not allowed to discriminate the premiums paid for that insurance according to individual criteria such as age, gender, and health status. Given that employer-provided insurance is the dominant source of insurance, it is a reasonable assumption for this question.

Thus, firms offer a full insurance at price ω , and agents select into the contract according to their demand. The firm's per-contract profit is:

$$\Pi = \omega - \mathbb{E}[\tilde{x}|\pi(\lambda) \geq \omega].$$

Free entry drives the expected profits of insurance firms to zero. Thus, the supply curve is the average cost of supplying insurance to the first $1 - \Gamma(\pi^{-1}(\omega))$ agents who select into it. The shape of supply curve is determined by the mechanics of marginal selection at every price. The supply curve may or may not be monotonic, depending upon the expected cost of the each agent, relative to the average cost of the agents up to that agent.

An equilibrium is a price of insurance, ω , and set of agents, $\Lambda \subset \Gamma$, who select, or pool, into that insurance contract, that are consistent with price equals average cost. Five equilibria types are possible:

- Downward-sloping demand and upward-sloping supply: in this case, the marginal agents are more expensive than the average agent;
- Downward-sloping demand and supply with, from left to right:
 - demand going from above to below supply;
 - demand going from below to above supply;
 - supply and demand just tangent, with supply above demand; and
 - supply and demand just tangent, with demand above supply.

⁵ It will be assumed that marginal utility does not change with the realization of the risk, besides the risk's effect on the agent's consumption level. If the risk changed the agent's marginal utility, then the agent may want consumption levels that are unequal between different outcomes of the realization. While such circumstances may be interesting, state-dependent utility functions are beyond the scope of study here.

The first equilibrium exhibits local favorable selection, as in Jovanovic [9]. This could be due to heterogeneity in risk preferences that was negatively correlated with expected cost. Alternatively, the final four equilibria exhibit local adverse selection. This typology of equilibria is necessary because different types of equilibria have different comparative statics. For the balance of the article, the tangency equilibria will be ignored, as they are not robust to changes in the supply and demand curves—either the equilibrium vanishes (supply and demand no longer cross), or two equilibria emerge from the changes (they further overlap).

Consider the second and third types. Suppose there is an outward shift of demand without a shift in supply—for example, if agents become more risk averse. When demand starts above supply, the new equilibrium supports more insurance at a lower price. The price falls because the selection is locally adverse. However, when supply begins above demand, the shift in the demand curve leads to less insurance at a higher price. These two cases are demonstrated in the sub-figures of Figure 1.

The response to a change in preferences can be characterized according to an equilibrium’s stability.

Definition 1. An equilibrium characterized by its marginal agent λ_m is locally stable, if, for all local $\varepsilon > 0$,

$$\omega(\lambda_m + \varepsilon) \geq \pi(\lambda_m + \varepsilon)$$

and

$$\omega(\lambda_m - \varepsilon) \leq \pi(\lambda_m - \varepsilon).$$

The stability refinement is similar to the trembling-hand refinement of Selten [10], and the stability refinement of Kohlberg and Mertens [11]. This refinement is intuitive—the equilibria in this model are pooling equilibria. The first and second equilibria are stable, while the third is not.




Figure 1. Supply and demand for insurance with asymmetric information and changes in demand.

Changes in preferences are generally difficult to identify, whereas it may be more feasible to identify changes to the distribution of risk, Γ . Changes to the distribution of risk can be classified as symmetric or asymmetric. First, consider asymmetric changes. Such changes have consequences for the average cost (supply) curve, with relatively modest effects on the marginal benefit (demand) curve. Figure 2 presents the three cases of interest—a (stable) equilibrium with favorable selection, and two equilibria exhibiting adverse selection, one stable and the other unstable.

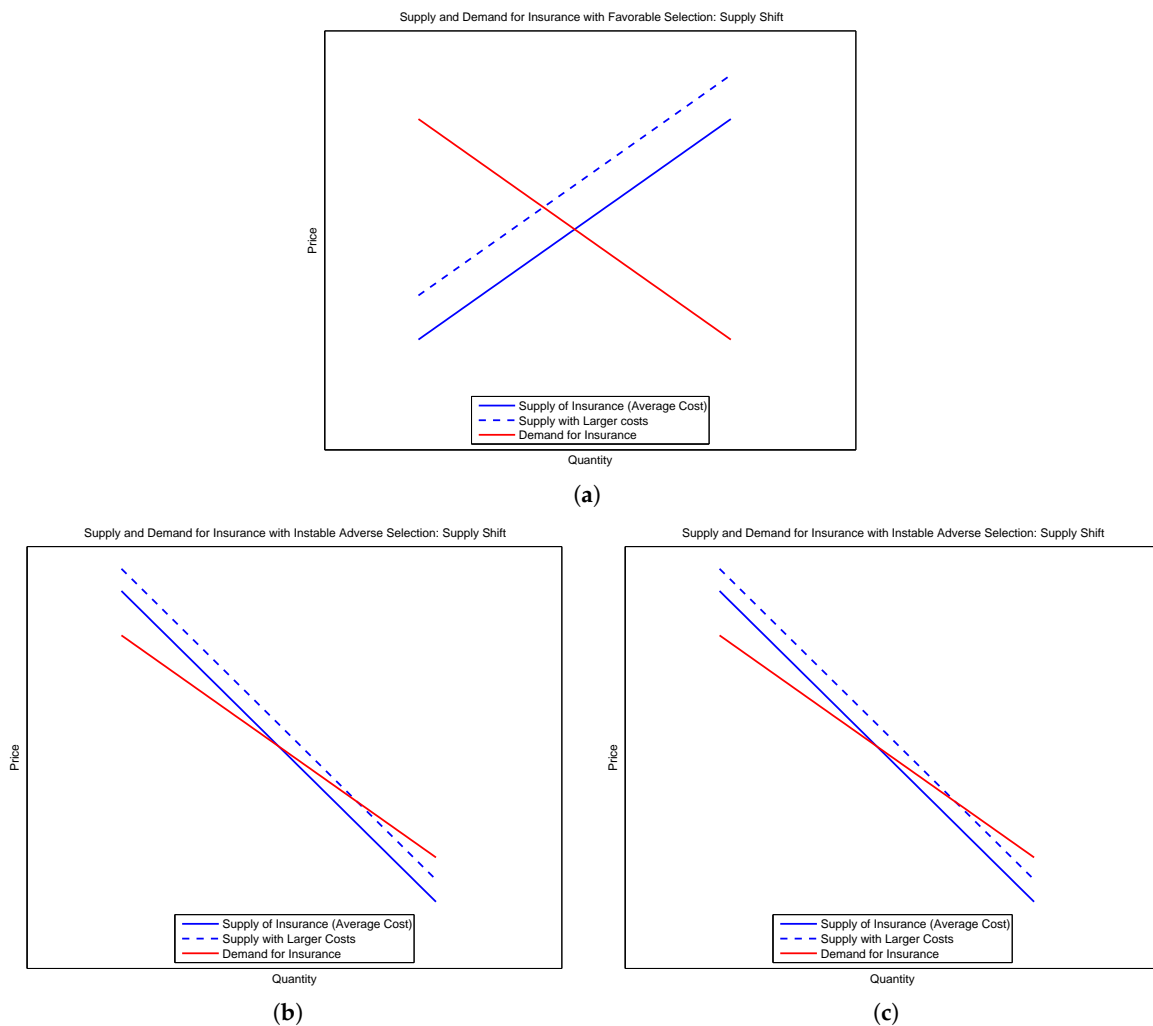


Figure 2. Supply and demand for insurance with asymmetric information and asymmetric changes to cost (i.e., no change in marginal demand).

The comparative statics depend upon the stability of the equilibrium. A shift up in the supply curve without a commensurate shift in the demand curve will lead to less insurance (a smaller pool) if we consider a stable equilibria. However, if we consider unstable equilibria, the opposite is true—a shift up in the supply curve will lead to an increase in the amount of insurance provided in spite of asymmetric information.

The number and stability of equilibria depend upon the micro-foundations of the model—preferences, and risk types and distributions, for example. In order to more clearly understand the relationship between risk and insurance in spite of asymmetric information, more assumptions are required.

3. A Series of Instructive Numerical Examples

Consider an economy with two types of agents who face a Bernoulli risk—they may catch an illness and have to go to the doctor. The cost of a visit, L , is the same for both types, though the healthy differ from the sick in that they face a smaller probability of going to the doctor, $p_h < p_s$. Fraction s of the population is sick, while the rest are healthy. The specific values for these parameters and

calculations for each example can be found in Table 1. If preferences for risk can be characterized by CARA, with degree of risk aversion r , each agent’s willingness to pay for full insurance is:

$$\pi(p., L; r) = \frac{1}{r} \ln \left(1 - p. + p.e^{rL} \right).$$

Note that $\pi(p., L; r)$ is an increasing function in each of its arguments. Sicker agents are willing to pay more for insurance, $\frac{\partial \pi(p., L; r)}{\partial p} > 0$, and all agents are willing to pay more for full insurance as the cost of a visit grows, $\frac{\partial \pi(p., L; r)}{\partial L} > 0$. It will also be useful to think of each agent’s certainty equivalent, $\pi(p., L; r) - p.L$ —the amount an agent is willing to pay for insurance above and beyond a fair price.

Table 1. Parameter values for Bernoulli numerical examples in the text.

Example	L	p_s	p_h	s	a_{all}^*	π_h	Adverse Selection?
1	1200	0.5	0.1	0.5	360	316	Yes
2	1500	0.5	0.1	0.5	450	483	No
3	1500	0.8	0.1	0.5	675	483	Yes

The average medical cost in the economy is $a_{all}^* = (sp_s + (1 - s)p_h)L$. Note that this value can increase because of increasing cost, L , or increasing probability of illness, $p.$. Suppose that the population is half sick and half healthy ($s = 0.5$). Furthermore, let $L = 1200$, and set the probabilities of illness to $p_s = 0.5$ and $p_h = 0.1$. The average medical cost is now 360, which would be the price of insurance if all agents chose full insurance bought in competitive markets with asymmetric information—the insurers do not provide type-based insurance. However, the healthy may not choose to buy such insurance, as their willingness to pay for insurance is less than its unfair price. This is the case here, using a calibration, $r = 0.0018$ —a healthy agent is only willing to pay 316 for insurance. The healthy go uninsured in such a market, while the sick sort into fairly-priced insurance.

Suppose that the price of a doctor’s visit rises to $L' = 1500$ —a symmetric change in risk. Average medical costs rise to 450. Now note, though, that the healthy agents are willing to pay for unfairly-priced full insurance provided by a competitive market. That is, $\pi(0.1, 1200; 0.002) = 483 > 450$. As the cost of a visit grows, the certainty equivalent of the healthy agent grows faster than the price of insurance. More full insurance was provided because these changes were symmetric—both the healthy and the sick face larger risks.

Now suppose an asymmetric change, $p'_s = 0.8$, keeping the price of a doctors visit at 1500. Average medical costs rise to 675. The healthy are not willing to pay this much for full insurance, as their willingness to pay for insurance does not change. Here, there is a decrease in insurance because the change was asymmetric.

This specification for risk is clearly not useful when considering the U.S. medical insurance market, as it is very coarse and does not allow for much heterogeneity. That said, it does suggest the general principle at work in the following section: following the average medical cost is not enough to determine whether an insurance market with incomplete information will allow for more insurance or less. The change in the distribution of risk that induced the change in the average medical cost is what determines the amount of insurance provided under asymmetric information. We will now apply that general principle to the preferred specification for medical risk.

4. A More General Setting

The model economy has a unit measure of risk-averse agents with utility function $u(\cdot)$. Each agent faces medical expenditures risk, $\tilde{m}x_i$, and its risk is private information. The risks faced by agents are heterogeneous and are distributed according to Γ . This model was previously used in Koch [12], which contains the details of preference parameter estimation. The private information assumption means that insurers are not able to sort, so all agents face a common price of insurance, ω . The purchase of

insurance is subsidized at a rate s , which is paid for using lump-sum taxes, τ . Insurance is subsidized at a rate here as it is in the U.S. federal income tax code—here it is a positive subsidy, relative to wage income. The two are equivalent with constant absolute risk aversion. Depending upon type, each agent has a willingness to pay for insurance, π_i , which makes them indifferent between facing the risk associated with their type, $\widetilde{m}x_i$, and paying for insurance. The decision to purchase insurance, ι_i , is a discrete choice:

$$U_I(w, \widetilde{m}x_i; \omega, \tau) = \max \left\{ u(w - (1 - s)\omega - \tau), E \left[u(w - \widetilde{m}x_i - \tau) \right] \right\}. \tag{1}$$

If $u(w - (1 - s)\omega - \tau) > E \left[u(w - \widetilde{m}x_i - \tau) \right]$ (i.e., $\pi_i \geq \omega$), then the agent purchases insurance, and does not otherwise. Because of asymmetric information, agents do not get access to fairly-priced insurance. The distribution of π_i is the demand curve for full insurance.

The market to provide full insurance is competitive, so that the price of insurance is equal to the average realized risk of the insured. Thus, we define an equilibrium to be:

- the insurance choice, ι_i , is the optimal solution to Equation (1);
- insurance price, ω , is equal to the average realized risk of the insured; and,
- tax level, τ , is equal to the total cost of the subsidy of medical insurance.

The specification of preferences and risk heterogeneity is chosen to mirror several empirical facts, while also aiding in the computation of equilibria. The cross-sectional distribution of medical expenditures has three main features: a large number of zeros, a monotonically-decreasing partial density, and a fat, Pareto-like tail.

In order to match the three main features of the distribution of medical expenditures, we will use the following specification of individual risk and heterogeneity:

- individual i faces risk according to the exponential distribution, parameterized by λ_i ; and,
- these types are distributed according to the Gamma distribution, with parameters α and β .

These two assumptions lead to a cross-sectional (i.e., unconditional) Pareto-type distribution of realized medical risk, with all three characteristics. This specification of uncertainty was originally proposed by Harris [13], though in a different context. Preferences against this risk are assumed to exhibit CARA.

CARA preferences are tractable, especially in this context. The willingness to pay for insurance, given type λ_i , parameter of risk aversion r , and upper bound of risk κ can be written as:

$$\pi(\lambda_i) = \frac{1}{r} \log \left(\frac{\lambda_i - r e^{-(\lambda_i - r)\kappa}}{\lambda_i - r} \right).$$

Note that willingness to pay for insurance with CARA preferences is independent of wages or wealth.

It can be shown that this willingness to pay is monotonically decreasing in λ_i . Because of this, equilibria can be characterized by their marginal agent type, λ_m , whose willingness to pay for insurance is equal to its price. The price of medical insurance is equal to the average realized risk of the insured:

$$\begin{aligned} \omega(\lambda_m) &= \int_0^{\lambda_m} t^{-1} \frac{t^\alpha e^{-\frac{t}{\beta}}}{\beta^{\alpha+1} \Gamma(\frac{\lambda_m}{\beta}, \alpha + 1)} dt \\ &= \frac{\Gamma(\frac{\lambda_m}{\beta}, \alpha)}{\beta \Gamma(\frac{\lambda_m}{\beta}, \alpha + 1)}, \end{aligned}$$

where $\Gamma(\cdot, \cdot)$ is the incomplete gamma function.

These specifications also mitigate the problem of multiple equilibria. Figure 3 plots out the supply and demand curves for a set of candidate parameters—the degree of risk aversion, r , upper bound of risk faced by an agent, κ , and the distribution of medical risk, α and β . Preferences and adverse selection (due to asymmetric information) lead to a downward-sloping supply curve. In this case, there are two potential equilibria. Other parameterizations could lead to one equilibrium (the supply and demand curve are just tangent), or no equilibria—the case where no risk sharing is possible.

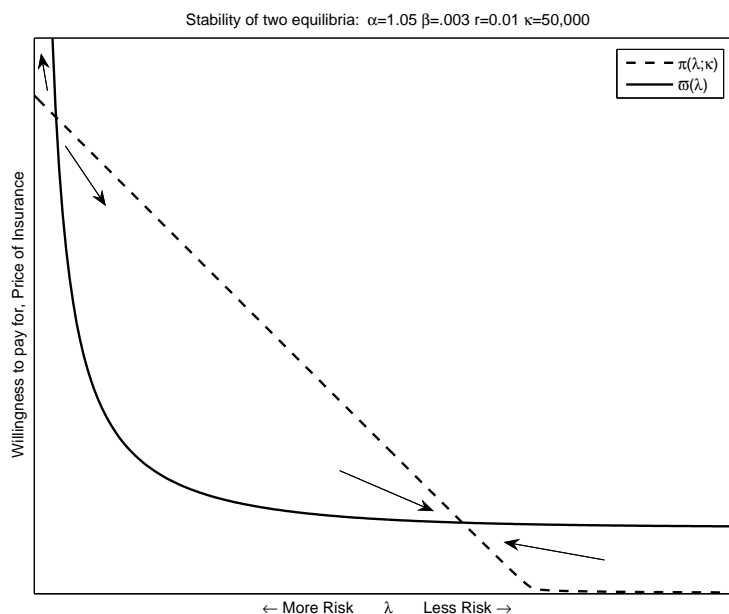


Figure 3. The stability of two equilibria.

Because the movements in insurance rates and prices are steady, it is unlikely that the equilibria we observe is unstable. Figure 3 demonstrates that the equilibrium with more insurance is the stable of the two. Parameter values that lead to a single equilibrium lead to unstable equilibria.

Medical insurance in the U.S. is primarily provided in the context of firms. Thus, the ideal data set would consider risk sharing across individuals within a firm. Unfortunately, such data is not publicly available. Instead, we use data from a nationally representative data set. To the extent that risks faced by workers within a firm look like the general population, this will be a useful examination. If workers could act on the distribution of medical risks within a firm, or firms could fire workers for being poor medical risks, this would be an unreasonable approximation. However, in both cases, such behavior might very well be either infeasible or illegal. Part of the medical risk under consideration here is latent—factors such as family history can impact future expectations—and information about it would thus be difficult to gather reliably. Even for evident medical conditions, it may be against the law for supervisors or co-workers to be informed about the health status of others, let alone act upon it.

Alternatively, one might worry that workers do not have an opportunity to take-up insurance within a firm. Decreases in the overall insurance rate may be due to fewer firms offering insurance, instead of fewer workers taking it up. This model does not account for these sorts of changes. Evidence suggests that these kinds of changes are not important. Stanton [14] finds that the number of workers employed by firms that offered medical insurance rose from 1996 to 2002. The decrease in employer-provided insurance over this period is instead tied to decreases in take-up by employees, and, to a lesser extent, fewer employees being eligible for medical insurance.

Here, we use the Medical Expenditure Panel Survey (MEPS). The MEPS is a series of overlapping two-year panels, with information on annual medical expenditures, charges, and insurance status. The panel aspect of the data is beyond the focus of this article—we use the series of cross sections in the

MEPS, starting in 1996, going to 2004, in order to infer each year's distribution of medical risk—each year's Gamma distribution, its (α_t, β_t) pair. We will later use these values to infer the amount of private insurance possible in the asymmetric-information model.

We restrict the sample to the non-elderly and the non-poor in order to avoid complications due to Medicare and Medicaid. Veterans are also removed, due to access to VA care. In all three cases, publicly-provided insurance can crowd out private insurance. This restriction also avoids potential confusion about private insurance among the elderly (over 65).

Unfortunately, the data do not present a perfect candidate for the insurable (i.e., covered by insurance) realizations of medical risk. One potential measure is the total charges for an individual over a year. Because medical insurance contracts typically have co-payments or deductibles, this is an imprecise measure of the realized insurable medical risk.

This specification of risk has an internally-coherent and tractable way to back out the insurable fraction of total risk. If a fraction $\rho^{-1} < 1$ of an exponential risk is insurable, parameterized by λ , then this fraction of a risk is also an exponential risk. This new risk is parameterized by $\frac{\lambda}{\rho}$. If this fraction ρ is common across agents, then all of the exponential risks are scaled by ρ . A common ρ might come from the fact that these co-insurance rates are used to solve a hidden-action problem, and type privacy does not allow for discrimination. Finally, since β is a scale parameter for the Gamma distribution, the new distribution of insurable risk is also a Gamma distribution, with parameters α and $\frac{\beta}{\rho}$.

The distribution of insurable risk can also be inferred from the distribution of medical expenditures paid by private insurance companies for the privately insured. The first moment of the conditional distribution, the average realized risk of the insured, can be found by integrating the expected realized risk over the types who choose insurance; i.e.,

$$\begin{aligned} E(\widetilde{m}x_i | t_i = 1) &= \int_0^{\lambda_m} t^{-1} \frac{t^\alpha e^{-\frac{t}{\beta}}}{\beta^{\alpha+1} \Gamma(\frac{\lambda_m}{\beta}, \alpha + 1)} dt \\ &= \frac{\Gamma(\frac{\lambda_m}{\beta}, \alpha)}{\beta \Gamma(\frac{\lambda_m}{\beta}, \alpha + 1)}. \end{aligned}$$

The average square of realized risk of the insured (i.e., the second non-central conditional moment) is found similarly:

$$\begin{aligned} E(\widetilde{m}x_i^2 | t_i = 1) &= \int_0^{\lambda_m} 2t^{-2} \frac{t^\alpha e^{-\frac{t}{\beta}}}{\beta^{\alpha+1} \Gamma(\frac{\lambda_m}{\beta}, \alpha + 1)} dt \\ &= \frac{2\Gamma(\frac{\lambda_m}{\beta}, \alpha - 1)}{\beta^2 \Gamma(\frac{\lambda_m}{\beta}, \alpha + 1)}. \end{aligned}$$

These are two unique moments that identify the two parameters of the risk distribution, (α, β) .

That said, the amount paid for treatment by a private insurance company is unlikely to be the amount billed to the uninsured. Insurance companies buy a large volume of medical goods and services throughout a year, and frequently bargain with with medical providers over the cost of treatment. This bargaining power is not available to the individual uninsured patient. The resolution here is the same as for charges—assume a common proportional mark-up from the expenditures paid for by private medical insurance companies. This new marked-up distribution of risk is also a Gamma distribution, with parameters α and $\frac{\beta}{\rho}$. The balance of this article will use both methods—marking up from private medical expenditures, and marking down from total charges. Both series tell the same story.

It may be useful to characterize changes in α and β . As the discussions of mark-up and mark-down above suggest, changes in β are easier to interpret. Since β is a scale parameter, a decrease in β scales up the medical risk faced by every agent an equal amount.

The parameter α is best characterized by a distribution's kurtosis. A distribution's kurtosis is the ratio of its fourth moment to the square of its second moment (variance). As discussed in Balanda and MacGillivray [15], kurtosis is synonymous with "peakedness" and "tail weight." Two distributions may have the same variance, but kurtosis measures whether that variance is due to a large amount of moderate dispersion or a moderate amount of large dispersion. Flatter distributions have lower kurtosis.

Changes in kurtosis are central to the ability of insurance markets to spread risk across types. The variance of risk types may increase. If this increase is because many types became a little more different, then it is unlikely that many will newly refuse to buy the same insurance. However, if a few of the agents became very different from the rest, then cross-type risk sharing becomes less likely.

The kurtosis for a gamma distribution is $\frac{6}{\alpha}$. Thus, as α decreases, the risk types are becoming more and more disparate, and market failure becomes more pervasive. The sickest members of the population are becoming relatively more sick than the rest of the population. We will refer to this effect as the *kurtosis effect*—as α increases, risk types become more similar, and insurance becomes more widespread. When κ is large, the kurtosis effect dominates.

However, the calibrated value of κ is only \$5600. This reflects the many ways in which uninsured individuals can seek and receive free or uncompensated care. Hospitals reported providing an average of \$20 billion per year in uncompensated and charity care (American Hospital Association [16]). The charity portion of care reflects the not-for-profit and humanitarian missions of many hospitals. Uncompensated care is typically provided in emergency rooms, where hospitals are required to provide care independent of insurance status or ability to pay. In either case, this provides an upper bound risk faced by an uninsured individual, and it has consequences for an agent's willingness to pay for insurance (i.e., in distribution, the demand curve for insurance). The calibrated value for CARA risk aversion, $r = 0.0018$, is in the range of risk preferences estimated by Cohen and Einav [17].

The subsidy rate, s , has important implications for the calibration. Because ninety percent of the privately insured get medical insurance from their employer, and the cost of that insurance is not taxed as wage income would be, the subsidy rate is the marginal tax rate faced by the individual. Marginal tax rates are important, because they provide extra incentive to have insurance, and increase the insurance rate. This model corresponds to a competitive labor market where workers are compensated with their productivity, and only their compensation can come in the form of taxed wages or untaxed medical insurance.

Saez [18] provides the average marginal tax rates for the US. These calculations stop at the year 2000. The marginal tax rate for the year 2001 and later will be the average value for the years 1996 to 2000.

When α decreases, the marginal agent faces a larger risk, and her willingness to pay for insurance increases. However, when that agent faces only a fraction of the realized risk (i.e., κ is small), the marginal agent's willingness to pay for insurance does not increase more than the increase in the price of insurance. Here, the price of insurance is increasing because everyone who has insurance is also becoming more expensive. The kurtosis effect is dominated by the price effect of decreasing α .

Recall the claim in Swartz [1] that increases in the average cost of care led to the increased price of insurance, and therefore fewer people had medical insurance. If the increase in average medical expenditure is due to a common increase in medical risk (β goes down), then the insurance rate may increase. As everyone's medical risk grows, so does their certainty equivalent—how much above and beyond a fair price each agent is willing to pay for insurance. Larger certainty equivalents allow for more risk sharing across types. If the increase in medical expenditure is the result of a change in the shape of the distribution of medical risk, then the outcome is ambiguous—either the price effect or the kurtosis effect may dominate.

As described in Table 2, the changes inferred from 1996 to 2004 are a mixture of the two. Over this period, we might derive these changes as coming from two sources: an increasing price of medical goods and services (falling β); and an increasing incidence of costly medical conditions (falling α). This

latter trend is consistent with the aging American population and the increased incidence of obesity, each associated with costly medical conditions.

The mark-up rate for a given year is exactly identified as the ratio of the price of insurance and the average medical expenditure paid for insurance on behalf of the insured, both from the data. The mark-down rate from total charges for a particular year is found by matching the model’s price of insurance to that found in the data. These estimates are presented in Table 2. The price of insurance used here is the average price of single-coverage employer-provided medical (hospital and doctor) insurance in the US, as reported in the MEPS. Since the risk is reported on the individual level, the price of the corresponding coverage is considered as well. As mentioned in the introduction, the price of single-coverage and family-coverage employer-provided insurance grew at similar rates. Over this period, family coverage (covering two or more individuals) was steadily two-and-a-half times as expensive as single coverage. The implicit assumption is that there are constant returns to scale when it comes to family insurance plans versus single coverage, and there is no differential selection across them.

Table 2. Parameter values that vary by year; see text for explanation of subsidy rate for years since 2000.

Year	Charges			Expenditures			Subsidy
	α_{year}	β_{year}	ρ_{year}	α_{year}	β_{year}	ρ_{year}	
1996	1.091	5.30×10^{-4}	0.9706	1.0453	0.0010	1.8046	0.2475
1997	1.107	5.17×10^{-4}	0.968	1.1557	0.0011	2.0931	0.2533
1998	1.087	4.70×10^{-4}	0.9079	1.1645	0.0011	2.1757	0.2556
1999	1.128	4.57×10^{-4}	0.988	1.1479	0.0010	2.1994	0.2584
2000	1.156	3.95×10^{-4}	1.0028	1.1203	0.0010	2.2978	0.2613
2001	1.132	3.51×10^{-4}	0.9145	1.166	0.0008	2.118	0.2522 *
2002	1.141	3.21×10^{-4}	0.9293	1.1838	0.0008	2.2149	0.2522 *
2003	1.049	3.01×10^{-4}	0.84	1.0547	0.0008	1.9889	0.2522 *
2004	1.067	2.72×10^{-4}	0.8126	1.1096	0.0007	2.0232	0.2522 *

Figure 4a,b plot the insurance rate predicted by the model against the rate observed in the data. The two trends are within four percentage points of one another until 2001, the year of significant tax reform legislation. Since this tax reform decreased the marginal tax rates faced by households, this may explain why the model predicts increasing insurance rates, while the data finds insurance rates dropping slightly since 2000.




Figure 4. Predicted trends are consistent across methods.

Figure 5a,b assess the relative importance of changes in the level of risk (β) and the shape of risk (α). They each plot two counterfactual insurance rates—the first, if the shape of the distribution of risk stayed the same, but the level of risk changed as described in Table 2 (i.e., α_{1996} for all years, β_t for all t); the second, if the level of risk did not change, but the shape of the distribution of risk changed as described in the same table (i.e., β_{1996} for all years, α_t for all t).




Figure 5. Changes in levels dominate changes in shape.

Changes common to the medical risk of all agents are driving the insurance rate trends predicted by the model. Inferred changes in the shape in the distribution of medical risk (i.e., changes in α) play a very small roll in the predicted trends. Thus, in this period, changes due to the shifting age distribution or obesity have very little explanatory power. Instead, the model suggests a focus on changes common to all agents, such as moving prices for medical goods and services.

5. Conclusions

A competitive insurance market's capacity to provide insurance in spite of asymmetric information depends upon the distribution of risk. When an economy faces greater average risk, the amount of uninsured risk may vary, but depends crucially upon what kind of change took place. An observed increase in average medical costs could be due to either a common increase in the risk faced by all; or, it could reflect the change in risk faced by some agents. In the former case, there is more opportunity for risk sharing across types. In the latter, uninsurance is more frequent because the risk types are more disparate.

Changes in the cross-sectional distribution of medical risk suggest that there have been changes to both the level and the shape in the distribution of medical risk. The changes in the medical insurance rate predicted here are due primarily to the former. This emphasizes the role of changing medical prices for the rate of and price for insurance. It meanwhile suggests that changes due to shifting demographics, such as age and obesity, have had little effect on the number of people with medical insurance.

Conflicts of Interest: The views expressed in this article are those of the author. They do not necessarily represent those of the Federal Trade Commission, and the work was not prepared as part of the author's work for the FTC.

References

1. Swartz, K. *Reinsuring Health: Why More Middle-Class People Are Uninsured and What Government Can Do*; Russell Sage Foundation Publications: New York, NY, USA, 2006.

2. Freudenheim, M. *Cost of Health Insurance Rises Again, but at a Slightly Slower Rate*; New York Times: New York, NY, USA, 2007.
3. Gruber, J.; Levy, H. The Evolution of Medical Spending Risk. *J. Econ. Perspect.* **2009**, *23*, 25–48.
4. Handel, B.R.; Hendel, I.; Whinston, M.D. Equilibria in Health Exchanges: Adverse Selection vs. Reclassification Risk. *Econometrica* **2015**, *83*, 1261–1313.
5. Koch, T.G. One pool to insure them all? Age, risk and the price (s) of medical insurance. *Int. J. Ind. Organ.* **2014b**, *35*, 1–11.
6. Charpentier, A.; le Maux, B. Natural catastrophe insurance: How should the government intervene? *J. Public Econ.* **2014**, *115*, 1–17.
7. Cummins, J.D. Should the government provide insurance for catastrophes. *Fed. Reserve Bank St. Louis Rev.* **2006**, *88*, 337–379.
8. Kousky, C.; Cooke, R. Explaining the failure to insure catastrophic risks. *Geneva Pap. Risk Insur. Issues Pract.* **2012**, *37*, 206–227.
9. Jovanovic, B. Favorable Selection with Asymmetric Information. *Q. J. Econ.* **1982**, *97*, 535–39.
10. Selten, R. Re-examination of the perfectness concept for equilibrium points in extensive games. *Int. J. Game Theory* **1975**, *4*, 25–55.
11. Kohlberg, E.; Mertens, J.-F. On the Strategic Stability of Equilibria. *Econometrica* **1986**, *54*, 1003–1037.
12. Koch, T.G. Bankruptcy, medical insurance, and a law with unintended consequences. *Health Econ.* **2014**, *23*, 1326–1339.
13. Harris, C.M. The Pareto Distribution as a Queue Service Discipline. *Oper. Res.* **1968**, *16*, 307–313.
14. Stanton, M.W. *Employer-Sponsored Health Insurance: Trends in Cost and Access*; Research in Action 17; Agency for Healthcare Research and Quality: Rockville, MD, USA, 2004.
15. Balanda, K.P.; MacGillivray, H.L. Kurtosis: A Critical Review. *Am. Stat.* **1988**, *42*, 111–119.
16. American Hospital Association. *Uncompensated Hospital Care Cost Fact Sheet*; American Hospital Association: Washington, DC, USA, 2014.
17. Cohen, A.; Einav, L. Estimating Risk Preferences from Deductible Choice. *Am. Econ. Rev.* **2007**, *97*, 745–788.
18. Saez, E. *Reported Incomes and Marginal Tax Rates, 1960–2000: Evidence and Policy Implications*; NBER Working Papers 10273; National Bureau of Economic Research, Inc.: Cambridge, MA, USA, 2004.



© 2017 by the author; licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>).