One Pool to Insure them All? Age, Risk and the Price(s) of Medical Insurance

Thomas G. Koch *

December 2011

Abstract

Asymmetric information can lead to adverse selection and market failure. In a dynamic setting, asymmetric information also limits potential future price variation. This certainty offsets the costs of adverse selection. Using a dynamic model of endogenous insurance choice and price calibrated to the U.S. medical insurance market, I find that asymmetric information is Pareto improving when information is fully asymmetric. However, when insurers can discriminate by age group, but not within age groups, the young benefit by paying less for insurance. The insurance market for the near elderly collapses because it is no longer implicitly subsidized by the participation of the young.

JEL Codes: D82, I1, L51

Keywords: asymmetric information, adverse selection, privacy laws.

*An earlier version of this paper was circulated under the title "Pareto Gains from Hiding Information." The views expressed are those of the authors and do not necessarily represent the views of the Federal Trade Commission or of any individual Commissioner. Federal Trade Commission, 600 Pennsylvania Avenue, NW, Mail Drop M-8059, Washington, DC 20580; Phone: 512-809-8014 E-mail: tkoch.at.ftc.gov.

1 Introduction

Asymmetric information can lead to adverse selection and market failure as risky "lemons" price the less risky out of non-discriminatory contracts. As market failure, this leads to squandered surplus—the uninsured lack insurance only because they lack access to fairly priced insurance. Thus, we get a common objection to information restrictions in the medical insurance market—in attempting to make the market "fair" (equal price for all), the market is not able to provide a set of "fair" (equal to expected cost) prices that would ensure full insurance. If US state and federal governments enact such restrictions, the welfare costs of adverse selection should be expected.

While adverse selection may lead to un-insurance, it also limits the variation of prices paid by those who do have insurance. To the extent that adverse selection creates type pooling, it creates price pooling. This price pooling mitigates future price risk.

This paper investigates potential welfare consequences of asymmetric information in a dynamic setting, where risk types (and, thus, fair prices) may change over time. I calibrate a dynamic model of insurance choice and risk pooling to the United States medical insurance market, and find that fully asymmetric information (i.e., one contract at one price) is Pareto preferred to fully symmetric information with fair pricing. However, when insurers can discriminate by age group, but not within age groups, the young benefit by paying less for insurance, while the insurance market for the near elderly collapses because it is no longer implicitly subsidized by the participation of the young. Thus, using information may only be a partial solution to adverse selection, because the cross-subsidization may be an important factor in some groups having insurance.

There is a general tension when evaluating the welfare consequences of asymmetric information. Asymmetric information may lead to complete pooling (everyone buys insurance), partial pooling (some buy insurance), or no pooling. With complete pooling, asymmetric information is not distortionary (everyone makes the same decision), but the new allocation is different in its distribution. With no pooling, asymmetric information is destructive, in that there are no contracts that satisfy participation constraints. (I.e., it is distortionary.) Without any insurance, there are no insurance prices to pool. With partial pooling, the market is only partially inhibited, so the policy is both distortionary and distributional.

These tradeoffs may inform analysis of recent and future policy changes. The federal Patient Protection and Affordable Care Act has provisions that will limit the amount of price discrimination for private insurance within age groups. The distortionary and distributional consequences of this reform are evaluated in a more general lifecycle model of insurance choices and prices. This model is calibrated to age-specific distributions of risk, and the ability of markets to provide insurance in spite of within-group asymmetric information. Partial information may offset the adverse selection in some groups (those that subsidize the high risk when no information is used in pricing), but lead to less insurance among those who value it the most (the high risk group that needs that subsidy to support the participation constraints with near-universal insurance).

This series of results can be interpreted as "second best" results. The first-best outcome shares the risk of realizing types, and having to pay for insurance accordingly. In the absence of this allocation, an equilibrium with some market failure may be preferable to one with perfectly symmetric information and many fair prices.

Hirshleifer (1971) associated the heterogeneity of risk types as a risk itself, faced by agents before types are revealed. Chiappori (2006) recast this risk as medical risk, where medical screening provides early revelations of relevant information about future health. In this context, there may be large welfare losses due to medical screening. The results of this paper are similar, but with a key distinction. Before medical screening is introduced, all agents are the same, and the welfare considered is that of the agent *ex ante* (before risk types are revealed by the screening). Here, the focus is on the welfare of all agents *ex interim*, each of whom knows their risk type. In a sensibly calibrated model asymmetric information is Pareto improving—all agents, sick and healthy, would prefer asymmetric information to symmetric information.

The uncertainty over medical risk type from one year to the next is identified by Cochrane (1995), which proposes a contract to insure against this risk. Here, asymmetric information provides insurance similar to that contract—by forbidding an insurer from discriminating according to risk type, there is insurance against becoming high risk (and having to pay a lot for insurance) in the future.

This result can potentially explain the myriad and persistent laws that enforce information privacy and their concomitant market failure. The employer-provided nature of medical insurance makes it hard for insurance companies to price discriminate; medical privacy laws further this end. If these practices are inefficient, then why do they persist? This result suggests that there are gains from non-discrimination. These are the gains found by Hendel and Lizzeri (2003) in their study of one-sided life insurance contracts—restricting the ability of life insurers to discriminate through one-sided commitment serves the same purpose as provisioning medical insurance through a third party (employer) who does not discriminate. This reclassification risk is uninsured in dynamic insurance markets where long-term insurance contracts are not enforceable.

This paper's result is consistent with the findings of de Garidel-Thoron (2005), which investigated asymmetries of information between an individual's current insurer and potential future insurers. In a two-period model with Bernoulli risks and monopolistic competition among insurers, asymmetries across insurers can induce more long-term insurance contracts, which provide insurance against dynamic uncertainty. In contrast, this paper's focus is the asymmetric information between individuals and insurers with competitive markets for insurance and a rich specification for risk heterogeneity.

Polborn, Hoy, and Sadanand (2006) studied the potential benefits of asymmetric information in the life insurance market. Comparing *ex ante* welfare, they find that information regulations in the life insurance market can be welfare improving when there is risk- and demand-type uncertainty. The risk in their study is a low-probability, high-cost risk associated with end-of-life costs. It provides a convincing case for information restrictions on the results of tests against such risk. This paper provides evidence for a broader restriction on the use of information in the medical insurance market.

This paper can also provide context to recent empirical work on asymmetric information and adverse selection in insurance markets. Finkelstein and Poterba (2006) found that insurance companies have access to information that might assist in determining more accurate risk profiles, but do not use it. This behavior is ascribed to "political economy" concerns. Here, I evaluate the efficiency of restrictions placed on information. This would be especially true for markets where a firm buying insurance for its employees tries to maximize the overall utility of its workers for a given price of insurance.

Finkelstein, Poterba, and Rothschild (2009) investigated the welfare consequences of non-discrimination in mandatory annuity markets. The consequences of non-discrimination are measured in terms of transfers (from the low risk to the high risk), as well as efficiency *ex interim*—once the types are realized. Because the market is mandatory, there is no welfare loss due to adverse selection. This paper includes this margin when calculating the efficiency consequences of non-discrimination, and uses them to compare these *ex interim* costs with the *ex ante* benefits of price non-discrimination.

2 Risk and Pooling Over Time

The following is an adaptation of a model developed by the author (Koch (2007)) in previous work. Here, this specification is designed to mimic aspects of the United States employer-provided medical insurance market, and extend it to include dynamic components. First, consider the one-period specification.

2.1 Pooling in one-period contracts

Consider an economy with a unit measure of heterogenous agents facing medical risk. Individual agents have CARA preferences, and face conditional risks according to a type-specific exponential distribution. That is, preferences can be characterized according to the utility function $u(w) = 1 - e^{-rw}$ with risk-aversion parameter r, while the probability that an agent must pay less than or equal to X for medical goods and services is:

$$P(x < X|\lambda) = 1 - e^{-\lambda X},$$

where λ is the agent's type. If the types (λs) are distributed according to the gamma distribution, parameterized by (α, β), then the unconditional distribution of realized medical costs is a Pareto-type distribution:

$$
\begin{aligned}
P(x < X) &= \int_0^\infty \left(1 - e^{-\lambda X}\right) d\Gamma(\lambda; \alpha, \beta) \\
&= 1 - \frac{1}{(\beta X + 1)^{\alpha+1}}.
\end{aligned}
$$

Pareto-type distributions are consistent with the cross-sectional distribution of medical expenditures—observations at zero, none below zero, a monotonically declining partial density, and a fat tail.

There are three potential equilibria: one under asymmetric information, a second under symmetric information, and a third with partial information. When information is fully symmetric, competitive insurance markets provide type-specific full insurance contracts to agents at a price equal to expected cost, λ^{-1}. When information is asymmetric, there may be partial pooling of risk types. Because insurers cannot price discriminate, one full insurance contract is offered. Agents adversely select into the contract, which is priced according to the expected cost of the agents who select into the contract. Risk aversion increases agents' willingness to pay for insurance beyond expected cost, and allows for the cross-subsidization of the high risk agents by the low risk agents inherent in the asymmetric information contract. The asymmetric information equilibrium is defined as:

- a decision rule, $\iota(\lambda)$, that has agents buy into the insurance contract when their willingness to pay for full insurance is greater than or equal to its price; and,

- a price of full insurance, ϖ that is equal to the average realized risk of those who choose insurance.

The third information setting, or regime, is partial information—insurers have access to information that is potentially correlated with risk, and offer insurance according to that information. The kind of information considered here is demographic, such as age and gender, that is relevant, verifiable, and known by all parties. Contracts may differ across groups, but there is asymmetric information within groups. Solving for the within-group contracts follows from the solution to a fully asymmetric information contract, mutatis mutandis.

Solving the asymmetric information economy requires a few steps. Agents with wealth w and Bernoulli utility function u face risk $\widetilde{x}$, and make the following insurance choice decision: buy insurance, $\iota = 1$ at price ϖ, or go uninsured, $\iota = 0$, and face the risk. This problem can be written as:

$$U_I(w, \widetilde{x}_i; \varpi) = \max_{\iota \in \{0,1\}} \left\{ E\Big[u(w - \widetilde{x})\Big], u(w - \varpi) \right\}. \quad (1)$$

Because of asymmetric information, agents can only choose to buy an unfairly priced full insurance contract for ϖ, or go uninsured.

With CARA preferences facing an exponential risk, an agent's willingness to pay for full insurance is:

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

where λ is the agent's risk type, r is the degree of risk aversion as before, and κ is the maximum payment (risk outcome) faced by the agent if uninsured. This last parameter relates to the practical realities of the market for medical care—because consumers frequently have access to free care (through the not-for-profit mission of some providers), or the the providers are required to provide it (in the case of emergency rooms), there is a maximum realized risk that agents face, and it has consequences for their demand for insurance. If an agent's willingness to pay for insurance is greater than price ϖ, then the agent will choose to buy unfairly priced insurance, and will not otherwise. The demand curve for full insurance is the distribution of $\pi(\lambda)$, according to the distribution of risk types, λs, given by the gamma distribution. Willingness to pay for insurance is a decreasing function of risk type, so it leads to a downward-sloping demand curve.

Because of adverse selection, the supply curve for insurance with asymmetric information is downward sloping, and there may be zero, one or two equilibria. The supply curve is the average realized risk of the insured. This can be calculated as follows:

$$\begin{aligned}\varpi(\lambda_m) &= \int_0^{\lambda_m} t^{-1} \frac{t^{\alpha} e^{-t/\beta}}{\beta^{\alpha+1}\Gamma(\lambda_m/\beta, \alpha+1)} dt \\ &= \frac{\Gamma(\lambda_m/\beta, \alpha)}{\beta\Gamma(\lambda_m/\beta, \alpha+1)},\end{aligned} \tag{2}$$

where $\Gamma(\cdot; \alpha, \beta)$ is the incomplete gamma function, and λ_m is the marginal agent choosing insurance. Because willingness to pay for full insurance is monotonically decreasing, the price depends upon the distribution of risk and the marginal agent's risk type, λ_m. An asymmetric information equilibrium in this setting is when the insurance choices of the agents are consistent with the average realized risk of those choosing insurance. Simply, it is where the supply curve and the demand curve intersect. It can be defined according to its marginal agent, whose willingness to pay for insurance is equal to its price. Concerns about hidden action and loading costs, and how they distort π and ϖ will be addressed in the Appendix.

Figure 1 plots a supply and demand curve for a single full insurance contract under completely asymmetric information, with two potential equilibria. Further equilibrium refinements point to the right-most equilibrium. As a high-insurance, low-cost equilibrium, it is a Pareto improvement over the left-most equilibrium. It also exhibits stability against agents deviating from their optimal insurance strategies, as the arrows in Figure 1 suggest.

2.2 Changes Over Time

The dynamic asymmetric information equilibrium is a repeated version of the static asymmetric information equilibrium. From this, value functions are calculated according to risk type and information regime. With an annual discount factor of δ, an agent's value function with asymmetric information depends upon the agent's type, and the amount of across-type risk sharing supported by the economy:

$$V_A(\lambda) \quad = \quad \sigma_A(\lambda) + \delta\left(p \cdot V_A(\lambda) + (1-p) \cdot \mathbb{E}_{\lambda'}[V_A(\lambda')]\right)$$

where

$$\sigma(\lambda) = \begin{cases} u(w - \varpi), & \text{if } \lambda < \lambda_m; \\ \mathbb{E}_{\widetilde{x}|\lambda}[u(w - \widetilde{x})], & \text{else.} \end{cases}$$

The amount of risk sharing across types (i.e., for which agents is $\lambda > \lambda_m$, and for which it is not) is determined by the same mechanism described in the static setting. The insurance rate and price of insurance used here are the same as determined in the static setting.

Alternatively, if information is symmetric, then agents can always get insurance priced at equal to the agent's expected cost.[1] In a competitive environment, there is no insurance against changing risk type, as in Hendel and Lizzeri (2003). The dy-

[1]In some extreme cases, because uninsured out-of-pocket medical spending is capped, while insurance payments are not, some agents may choose not to buy fairly priced insurance. E.g., in extremis, if $\kappa = 0$, then $\pi(\lambda) = 0$ for all risk types. This circumstance is accounted for in the simulations described below.

namic symmetric information equilibrium is a repeated form of the static symmetric information equilibrium. The value function for an individual with current risk type λ and symmetric information is:

$$V_S(\lambda) \quad = \quad u(w - \lambda^{-1}) + \delta(p \cdot V_S(\lambda) + (1-p) \cdot \mathbb{E}_{\lambda'}[V_S(\lambda')]).$$

The present day valuations associated with the partial information regimes will be forward looking—my present day value of asymmetric information includes both my current value, as well as my projected future value. To incorporate this, we start with the present day valuation V_{wA}^G for age group G, under the presumption that the individual stays within that age group indefinitely. This corresponds to the value of asymmetric information constructed above, only with age-group G superscripts to denote the different risk groups.

$$V_{wA}^G(\lambda) \quad = \quad \sigma_A^G(\lambda) + \delta_G \left(p^G \cdot V_A^G(\lambda) + (1-p^G) \cdot \mathbb{E}_{\lambda'|G}[V_A^G(\lambda')] \right)$$

where

$$\sigma(\lambda) = \begin{cases} u(w - \varpi^G), & \text{if } \lambda < \lambda_m^G; \\ \mathbb{E}_{\widetilde{x}|\lambda}[u(w - \widetilde{x})], & \text{else.} \end{cases}$$

There is a natural relationship for the within-group value function that ignores the transitions between age groups, and the value function, referred to as telescoping (subscript t) here, which accounts for it. For example, imagine an infinitely lived individual is young for two years and then old the rest of his or her life. The value function of two years of young spending is $(1-\beta^2)V(\lambda; young)$. The telescoping value function, V_t, that accounts for two years of youth and then an infinitely-lived old age is

$$V_t(\lambda; young) = (1 - \beta^2)V_{wA}(\lambda; young) + \beta^2 V_{wA}(\lambda; old).$$

These functions would not be valid in their standard use in dynamic macroeconomics. In that setting, individuals are making choices, and changes to futures returns might alter current investment decisions. Here, we are just graphing out the relationship between exogenous risk type and the value associated with information regimes, so there are no such consequences for these value functions.

In constructing these age-based value functions, there are three differences across age groups: the cross-sectional risk sharing possible within each group (based on age-specific gamma distributions of risk, and the same mark-up across groups), the risk-type transition pattern, and survival probability. The first two are easily obtained from the MEPS and are calibrated independently in the fashion as before, while the survival probabilities are taken from mortality tables made available by the U.S. National Vital Statistics System. The mortality tables provide the mortality rates for ten-year age groups, starting at 25. There are ten-year age groups starting at age 25, and a child group (with no presumed mortality) which includes those from age zero to 24. Once age 65, an individual leaves the private insurance market (and the model more generally) for Medicare. The mortality rate is included by adjusting the discount rate, $\delta_G = \delta * (1 - m_{age})$.

When agents transition from one age group to the next and do not draw a new risk type, they maintain their relative position in the new cdf, not the absolute value of their λs. That is, a 34-year old who has the median risk type of the 25-34 age band, enters into the next age band and does not have to draw a new risk type, her new risk type is the median for that next age 35-44 age group. If she does have to draw a new risk type, she draws it from the new distribution.

3 Results

3.1 Symmetric Versus Asymmetric Information

The first-best outcome for *ex ante* agents provides insurance against both realized medical risk and realized medical-risk type. Welfare is measured relative to this first-best outcome—how much would an agent be willing to pay to get the first-best equilibrium. Equivalent variations vary by risk type, because risk type determines either the price of insurance if information is symmetric, or whether or not the agent purchases insurance if it is not. (Because CARA preferences do not exhibit wealth effects, compensating variation and equivalent variation are the same.)

The parameter values can be found in Table 1. The two cross-sectional parameters of the risk distribution are estimated from the mean and variance of medical spending among the insured population, as described in the Appendix. The parameters not estimated directly here are the degree of risk aversion, r, and the cap on realized medical risk, κ. Both were estimated in previous efforts of the author (Koch (2007)). The coefficient of absolute risk aversion used here is similar to recent estimates of Cohen and Einav (2007), while the cap on realized risk is consistent with Mahoney (2010). The default threshold parameter reflects the relatively generous provision of free care in the US, as well as the strict regulation of emergency rooms and their mandate to provide care independent of insurance status or ability to pay (in spite of the potential for subsequent default). This calibration leads to eighty-four percent ($\bar{\iota}$) of the population paying \$3,468 (ϖ) for insurance, while the balance go uninsured.[2]

Table 1 also lists the two calibrated parameters for the dynamic components of

[2]In the MEPS, approximately 76 percent of the sample are privately insured during the course of a year. Because the incidence of insurance is not used in this calibration, the model does not match the insurance rates in the age groups considered below. When we discuss insurance across age groups, the model does not consistently over- or underpredict the amount of adverse selection.

the model. The first is the probability of type transition, p, calibrated as described in the appendix. The second is the annual discount factor, δ, which is set to 0.96.

Figure 2 plots the difference between an agent's compensating variation for asymmetric versus symmetric information, according to its risk type, as a fraction of 1999 US consumption. That is, how much consumption would an agent be willing to give up in order to go from symmetric to asymmetric information.

Note that the equivalent variations are all strictly greater than zero—in this calibration, a switch from symmetric to asymmetric information is Pareto improving, with compensating variances on the order of 3 to 4.5 percent of consumption. This may seem large, but in this calibration, the most expensive insurance contracts can cost upwards of $10,000. The welfare costs associated with uncertain insurance prices are large, compared to the welfare loss due to adverse selection. There is a second-best trade-off. With asymmetric information, agents are insured against becoming high risk and paying a lot for insurance; this insurance is afforded by a subsidy from lower risk agents who pay more for insurance than is fair. This cross-risk type subsidization means that the lowest-risk agents may go uninsured, since the single-period insurance contract is more expensive than their willingness to pay for insurance. However, this tradeoff provides (type) insurance when the agent needs insurance the most, at the cost of (realization) insurance when insurance is the least valuable. This tradeoff is worth it, even for those who currently go without insurance due to adverse selection.

By construction, an agent's equivalent variation may be negative—the agent is worse off with the switch from symmetric to asymmetric information. In the extreme case where $p = 1$ (today's risk type is the same as tomorrow's risk type), a low-risk agent who can buy inexpensive insurance with symmetric information is worse off, because that insurance is no longer available with asymmetric information.

With a single risk pool, there is enough type switching to generate a broad support of asymmetric information. The shape of the graph can be instructive in understanding this support. Those who benefit the most from asymmetric information are the highest risk individuals, who receive the largest subsidy by paying the single, non-discriminatory price. As we move along the graph to the right, the size of this subsidy falls, eventually crossing smoothly through the point where agents going from pay too little to too much for unfairly priced insurance. The other feature of the graph is where it starts to increase. The upward slope begins with the agents who are presently adversely selected from insurance. The slope is upward because the surplus lost to information restrictions falls as agents face less medical risk.

3.2 Life Cycle Comparisons

Life-cycle components are an important aspect of medical spending heterogeneity and insurance choice. In fact, in order to avoid the kinds of adverse discrimination described above, PPACA instituted the use of age-bands, within which price discrimination would be limited, but across which was allowed, so as to mitigate adverse selection. What impact do the life-cycle components have on the optimal use of information? Will differences in the risk distributions by age group lead to different amounts of risk sharing by age group?

Table 1 presents the data moments for each of the five groups, and the corresponding parameter calibrations. The last three columns of the table report the results of the calibration exercise: the incidence of insurance ($\bar{\iota}$) and its price (ϖ) within each age group. The table also reports which fraction of the age group would have insurance ($\widetilde{\iota}$) in the market with fully asymmetric information. Not surprisingly, the amount of insurance among the young grows as they get age-specific insurance.

This is because they no longer implicitly subsidize the old in the non-discriminatory contract (paying 3,468 versus 1,800 for insurance for the youngest group). This result holds for the three youngest age groups.[3]

Markets are not able to pool risk for the older agents, and the markets for the oldest two groups collapse due to adverse selection. The reason they had insurance when information was fully asymmetric was because they were being subsidized by the young. The markets fail to work without that subsidy.[4]

Figure 3 presents the value functions at the start of each age group (18, 25, 35, etc., years). All individuals in the two youngest age groups prefer age-based pricing to either symmetric or fully asymmetric information. The former information regime provides them with a certain price of insurance, without subsidizing the older groups, as they would have to do in the final information regime. These gains are offset by the future (but discounted) adverse consequences of the age-based pricing when they get older.

The adverse consequences of age-based insurance pricing are most evident with the oldest group. When agents start at age 55, symmetric information is the same as age-based pricing, because they do not have insurance in either case. This is because the cap on medical expenses (calibrated to be $5,600) limits the demand for insurance to the point that they would not want to pay a fair price for it. This

[3]The model does not accurately predict the amount of insurance within these age groups. For example, the model predicts that 40 percent of the youngest group would have insurance with fully asymmetric information, while roughly 63 percent of them do in the data. Alternatively, the model underpredicts the insurance rate with the older groups—the two oldest groups have predicted insurance rate in the 90 percent range, while they are between 76 and 78 percent in the data. The general trends by age group—insurance rates go up as age goes up—is preserved. The lack of a consistent pattern suggests that the results that follow are not systematically over- or understating the implications of age-based pricing.

[4]These results were graphically verified—when the supply and the demand curves were plotted out across a dense grid of CDF values, the demand curve was below the supply curve at all points. The equilibrium values (and their uniqueness) for the other parameter values were confirmed in this manner, as well.

mechanism has little effect on the previous analysis, because such agents are rarely the marginal agent. Even in this extreme population, sixty percent of agents have a willingness to pay for insurance that is larger than an actuarially fair price. (For the two youngest groups, it is under ten percent.) However, the lack of a subsidy from the young drives these people out of insurance, and leaves them uninsured because of the default option (symmetric information), or never insured because an incentive-compatible risk pool cannot be constructed against the asymmetry of information. Thus, the near elderly are all made weakly worse off with age-based pricing. Only those who do not value insurance enough to pay a fair price for it are indifferent between symmetric information and age-based pricing. They are exactly indifferent because there is no chance of switching risk types in this age group.

The slopes of the two compensating variations are different. When age-based pricing is compared to the fully asymmetric regime, the slope is determined by the change in total willingness to pay for insurance. Thus, as agents become less risky, the loss of insurance associated with age-based pricing is less damaging. (All pay the same price for insurance with asymmetric information.) When age-based pricing is compared to the fully symmetric regime, the slope is determined by consumer surplus for insurance—willingness to pay less the price paid for insurance. The lowest risk agents are adversely selected out of insurance when information is fully asymmetric, so they are indifferent between the two. As agents become less risky, willingness to pay for insurance falls, but so does the price for insurance with fully symmetric information. For all but the lowest risk agent, the latter falls faster than the former, and the compensating equivalents for symmetric information grow as agents become less risky. The two graphs cross where the agent's expected cost (and, thus, fair price for insurance) is equal to its price under a fully asymmetric information regime.

These effects are evident, though less clear in the next youngest age group (starting at 45). Age-based pricing fairs worse than the other two information regimes, for similar reasons. Only 20 percent of this age group prefers uninsurance to a fair price for it, so there is some uninsurance in the symmetric information regime. These agents are still strictly worse off with symmetric information, because they may change risk type and lose insurance under age-based pricing. Comparing age-based pricing to symmetric information does not lead to a monotonic relationship between compensating variation and risk type, because the consumer surplus from a fairly-priced contract with symmetric information falls at roughly the sixtieth percentile. Monotonicity when comparing fully asymmetric information to age-based pricing is evident, as it is with the older age group.

The savings to the young in age-based pricing are evident in the plot for their age group—age-based pricing is preferred to fully asymmetric information by all in the two youngest groups. The discrete jumps in the compensating variation correspond to gains in insurance across the regimes. E.g., the jump in certainty equivalent at the fortieth percentile is due to the fact that such an agent is the marginal agent when insurance contracts are priced under the fully asymmetric information regime. Other jumps coincide with such gains and losses in insurance in the future; because of the way that risk types shift across age groups, the percentiles with jumps in the graphs do not map to the insurance rates in Table 1. The modest slope of the compensating variations is due to the future, steeper relationship in comparing age-based pricing to no information pricing, as described in the previous paragraphs.[5]

Most, though not all, of the youngest group does prefer age-based pricing to asymmetric information. If the young were always young, then the switch from

[5]Though not included in the welfare analysis here, the cost of financing the defaulted medical costs, if distributed evenly across the age groups, would offset some of the welfare gains of age-based pricing for the young.

symmetric prices to within-group asymmetric information would be Pareto improving within the group. This is because this group has such a low autocorrelation of medical spending, which suggests a high probability of switching risk types. This is also true of the next two age groups, starting at 25 and 35 respectively. However, these present day gains are offset by the fact that age-base pricing will be destructive when they enter the age groups starting at 45.

These results rely upon the calibration described in Table 1. Different calibrations may lead to different results. In particular, it may be interesting to know if other calibrations might lead to partially functioning insurance markets among the age groups, instead of their complete collapse.

These robustness checks can be classified into two different categories: adjustments to the preference parameters, and adjustments to the risk distribution. Insurance within the oldest age groups might persist in spite of the policy change if the simulations systematically undervalued insurance. For example, if the true calibrations of risk aversion were larger (i.e., increase r and κ). For example, doubling the degree of CARA risk aversion from 0.002 to 0.004 would lead to insurance in the second-to-last age group with age-based pricing, though not in the oldest age group. However, this would imply insurance rates well above 95 percent for all risk groups. Simulations where the cap on medical risk is moved up to $\kappa = 10,000$ have insurance in all age groups, but the insurance rates are likewise much larger than in the data. Increasing $\kappa = 10,000$ while decreasing $r = 0.001$ does keep the insurance rate with fully asymmetric information to 83 percent; insurance exists within all age groups when insurance contracts are priced by age group. However, the price of insurance under fully asymmetric information in this simulation is nearly $4,000, which limits the exercises's credibility. A medical risk cap this large would also be inconsistent with the findings of Mahoney (2010).

Alternatively, the calibration for the distribution of risk types could be wrong. Simulations with alternate distributions of medical risk were also considered. In particular, mark-up of medical risk is adjusted, plus and minus eleven percent. These are changes to the second parameter of the Gamma distributions, for all age groups. Because the second parameter of the Gamma distribution is a scale parameter, this scales up (down) all agents risk by 11 percent. (Why this is a "mark-up" is described in the Appendix.)

These simulations had similar results—changes to the scale of medical risk could make within-group insurance feasible, but not without other problematic results. Changes to the risk distribution shift both the supply and demand curves for insurance in the same direction, so it should not be any surprise that such changes do not overturn the qualitative results of the preferred calibration.

One general pattern was consistent across robustness exercises. Switching from fully asymmetric information, to asymmetric information within age groups, creates winners and losers. The youngest age groups gain in two ways—they are able to support larger risk pool, and at a lower price. The older age groups lose—they lose the implied subsidy from the young when age is not used in the pricing of insurance contracts. This increases the price for age-specific insurance contracts, which drives some, if not all, out of the insurance market.

4 Conclusion

Asymmetric information can lead to adverse selection and market failure as risky "lemons" price the less risky out of non-discriminatory contracts. This is the classical welfare cost of asymmetric information. However, the potential gains due to non-discriminatory pricing are large. The gains depend upon the amount of price-risk sharing induced by asymmetric information. If the pooling types are similar, the welfare gains may not be greater than its costs in un-insurance.

These two information structures are the opposite ends of a spectrum. This paper compares the *ex ante* welfare of those two end points. It would be interesting to know whether welfare is monotonically related to the usefulness of an imperfect measure of type. Similarly, the preferences of an *ex ante* agent map imperfectly to the circumstances faced by agents in a dynamic setting. If risk types are i.i.d, or sufficiently close, then these results may hold. However, as today's type more closely determines tomorrow's, the Pareto benefits of asymmetric information vanish.

This result is taken to a more general lifecycle model of insurance choices and prices. In particular, the lifecycle model considers age- (but not health-) specific pricing. This exercise highlights the large cross-subsidies involved in insurance markets that do not price according to relevant demographic information. Partial information may offset the adverse selection in some groups (those that subsidize the high risk when no information is used in pricing), but lead to less insurance among those who value it the most (the high risk).

These results highlight the complicated relationship between information and welfare. They demonstrate that partial information is only a partial solution to the risk pooling problem, and may make some groups worse off as a result. Policies that aim to limit market failure for some groups may aggravate the tenuousness of markets for other, related markets.

This is, of course, just one exercise. This exploration focused on age, but related exercises could look at gender, health or behavioral characteristics. While the tangled web of winners and losers for each policy may change, the interconnectedness of these results will likely remain.

References

Chiappori, P.-A. (2006): *Insurance: Theoretical Analysis and Policy Implications*chap. The Welfare Efects of Predictive Medicine. MIT Press.

Cochrane, J. H. (1995): "Time-Consistent Health Insurance," *Journal of Political Economy*, 103(3), 445–73.

Cohen, A., and L. Einav (2007): "Estimating Risk Preferences from Deductible Choice," *American Economic Review*, 97(3), 745–788.

de Garidel-Thoron, T. (2005): "Welfare-Improving Asymmetric Information in Dynamic Insurance Markets," *Journal of Political Economy*, 113(1), 121–150.

Finkelstein, A., and J. Poterba (2006): "Testing for Adverse Selection with Unused Observables," Working Paper 12112, National Bureau of Economic Research.

Finkelstein, A., J. Poterba, and C. Rothschild (2009): "Redistribution by insurance market regulation: Analyzing a ban on gender-based retirement annuities," *Journal of Financial Economics*, 91(1), 38–58.

Hendel, I., and A. Lizzeri (2003): "The Role Of Commitment In Dynamic Contracts: Evidence From Life Insurance," *The Quarterly Journal of Economics*, 118(1), 299–327.

Hirshleifer, J. (1971): "The Private and Social Value of Information and the Reward to Inventive Activity," *American Economic Review*, 61(4), 561–574.

Koch, T. G. (2007): "Bankruptcy, Medical Insurance and a Law with Unintended Consequences," Unpublished manuscript.

MAHONEY, N. (2010): "Bankruptcy as Implicit Health Insurance," Mimeo, Stanford University.

POLBORN, M. K., M. HOY, AND A. SADANAND (2006): "Advantageous Effects of Regulatory Adverse Selection in the Life Insurance Market," *Economic Journal*, 116(508), 327–354.

ROTHSCHILD, M., AND J. E. STIGLITZ (1976): "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," *The Quarterly Journal of Economics*, 90(4), 630–49.

A Estimating the Cross Section of Risk

The underlying gamma distribution can be fit to the ex-post unconditional distribution of realized risk, using a method of moments estimator. The two moments are the mean and variance of medical expenditures. The mean medical expenditure in the economy is:

$$\begin{aligned}\overline{x} &= \int_0^\infty \lambda^{-1} d\Gamma(\lambda; \alpha, \beta) \\ &= \frac{1}{\alpha\beta},\end{aligned}$$

while the variance is given by:

$$\mathrm{Var}(\widetilde{x}) = \frac{\alpha + 1}{\alpha^2 \beta^2 (\alpha - 1)}.$$

Unfortunately, the data do not present a perfect candidate for insurable medical risk. One potential measure is the total expenditures for an individual over a year. However negotiation between insurers and care providers might lead to an understatement of risk. If an insurance company gets a group discount, or mark down, from a provider, using an insurer's payment to a provider understates the amount the same individual would pay if he or she were uninsured.

This specification of risk has an internally-coherent and tractable way to back out a marked up distribution of risk. If there is a mark down rate for a private insurer of $\rho^{-1} < 1$ of an exponential risk, 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 mark downs provided to insurance companies are common, 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 $\frac{\beta}{\rho}$.

These mark downs also correspond to the multiplicative factors used to inflate the price of insurance in ways that correspond to loading costs and hidden action. The source of the mark down is not material here, so long as the prices and and willingnesses to pay for insurance against pure, exogenous risk employed here can account for any wedge that might be placed between the two.

The distribution of insurable risk can also be inferred from the distribution of medical expenditures paid for by 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.,

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

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

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

Such a fit to the cross-sectional distribution of medical expenditures can be found in Figures 4 (for the pdf) and 5 (for the cdf). The calibration in Table 1 is fit to these cross-sectional moments, as found in the Medical Expenditure Panel Survey (MEPS) for 2001-5. The MEPS is a series of nationally-representative panels, that measures medical care, utilization, expenditures by source and use, and insurance status. These conditional means are calculated using the weights provided in the data to account for the complex sampling used in the construction of the panels. The calibration has been adjusted to reflect the potential mark-up due to insurer-provider negotiation. This mark-up rate is set to match the average price of medical insurance in the middle of the sample (2003) and overall insurance rate, also provided by the MEPS.

B Calibrating Type Switching

How does today's risk type inform expectations about tomorrow's risk type? Two extremes can be informative. If risk types do not change over time, then there are no price-smoothing gains from asymmetric information. However, if tomorrow's risk type is an unconditional draw from the distribution of types, the welfare gains of asymmetric information could be large, on par with those measured previously, after appropriate discounting.

In general, today's risk type sheds *some* light on tomorrow's type, but may not completely reveal it. In equilibrium, the relationship between today's type and tomorrow's type must be consistent with tomorrow's distribution. That is, the transition function, $H(\lambda'|\lambda)$, for a steady-state equilibrium must satisfy:

$$\int H(\lambda'|\lambda) d\Gamma(\lambda; \alpha, \beta) = \Gamma(\lambda'; \alpha, \beta).$$

A Bernoulli draw, independent of risk type, is one such transition function. With probability p, an agent keeps its risk type, λ. Otherwise, the agent draws its risk type for next period from the original gamma distribution. The calibration of p has dramatic consequences for the welfare comparisons—if p is close to one, then the adverse selection costs of asymmetric information dominate; as it goes to zero, the agent grows more similar to the *ex ante* agent.

The values of p are calibrated using a simulation method. For each candidate p, 10,000,000 draws are taken from the fitted Gamma distribution, to determine first-period types. A second 10,000,000 Bernoulli draws are taken according to p to determine whether or not each agent switches types, and a third 10,000,000 draw for the (potential) second draws. The correlation coefficient among those agents who would select insurance in both periods is calculated for each p. Because the relationship between correlation coefficients and p may be noisy due to the randomness of the draws, the ps are linearly regressed on the correlation coefficients. The fitted values from that regression are used to match a correlation coefficient from the data to a probability of switching.

The Bernoulli probability p is calibrated to match the cross-sectional autocorrelation between insurance-paid medical expenditures in 2001 and 2005. If $p = 0$, then these observations are two random variables drawn from different distributions, and are thus uncorrelated. As p increases from zero, their correlation grows. Even if $p = 1$, they are not perfectly correlated, since they are just draws from the same distribution. This correlation is 0.306 in the MEPS, implying a transition probability of $p = 0.72$.

Table 1: Calibration for the dynamic model

			Preference Parameters, All ages					
		r	κ	λ_{max}	δ			
		0.002	5,600	2×10^{-4}	0.96			
			Risk Distribution					
Data Moments				Parameters		Model prediction		
$E[mx_i\|\iota=1]$	$E[mx_i^2\|\iota=1]$	Correlation	α	β	p	ϖ	$\bar{\iota}$	$\widetilde{\iota}$
		All Ages ($\delta = 0.96$)						
2066.36	6.38E+07	0.306	1.13	2.40E-04	0.72	3,468.79	0.844	–
		18-24 ($\delta = 0.96$)						
990.97	4.60E+07	0.063	1.04	6.12E-05	0.12	1,801.49	0.740	0.391
		25-34 ($\delta = 0.959$)						
1492.20	3.54E+07	0.074	1.12	3.47E-04	0.12	2,540.61	0.861	0.668
		35-44 ($\delta = 0.958$)						
1671.18	3.50E+07	0.201	1.16	2.82E-04	0.42	2,931.72	0.877	0.768
		45-54 ($\delta = 0.956$)						
2389.33	5.42E+07	0.301	1.23	1.83E-04	0.7	—	0	0.925
		55-64 ($\delta = 0.951$)						
3703.99	1.67E+08	0.524	1.17	1.25E-04	1	—	0	0.988

Estimated moments from the data (MEPS 2001-6) are those used in the exercises described in the text. The correlation presented is between private-insurance medical spending in both years of observation (e.g. , 2001 and 2002), among those who had insurance both years. The parameter values are calibrated directly from these values in the data. The three model predictions describe the difference between the market with and without age discrimination—ϖ and $\bar{\iota}$ are the price of and fraction with insurance when there are age group-specific insurance markets. $\widetilde{\iota}$ is the fraction with insurance within each age group when insurance markets do not discriminate by any characteristics. Differences between the average private insurance spending and the price of insurance in groups where everyone has insurance are due to (1)the mark-up as described in the text, and (2) the high-risk types moved from the distribution ($\lambda < \lambda_{max}$) correspond to different fractions of these different groups.

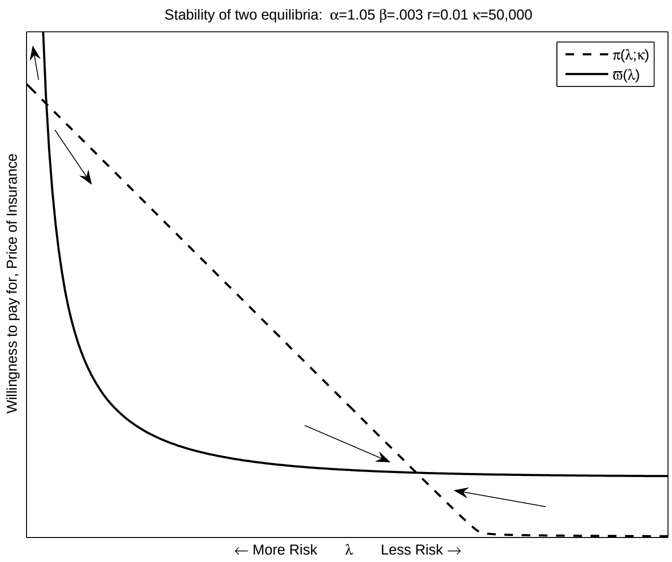

Figure 1: The equilibrium with a lower insurance rate is not stable, while the equilibrium with the higher insurance rate is stable.

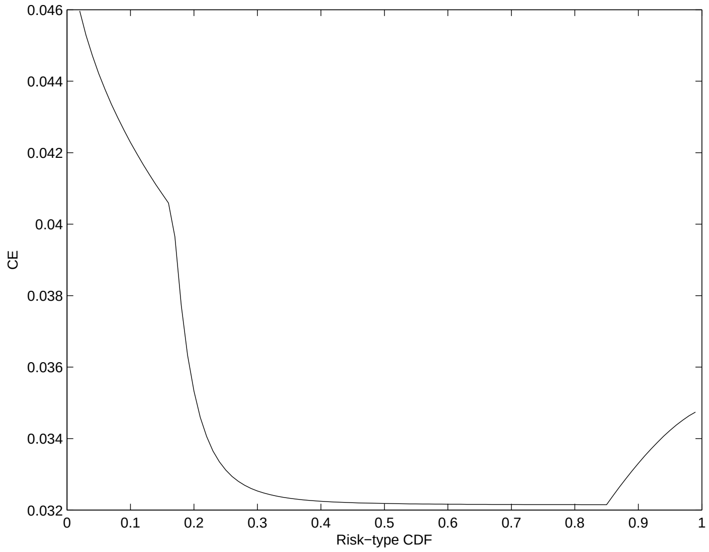

Figure 2: Compensating equivalences for symmetric-to-asymmetric information, with standard calibration for type transition.

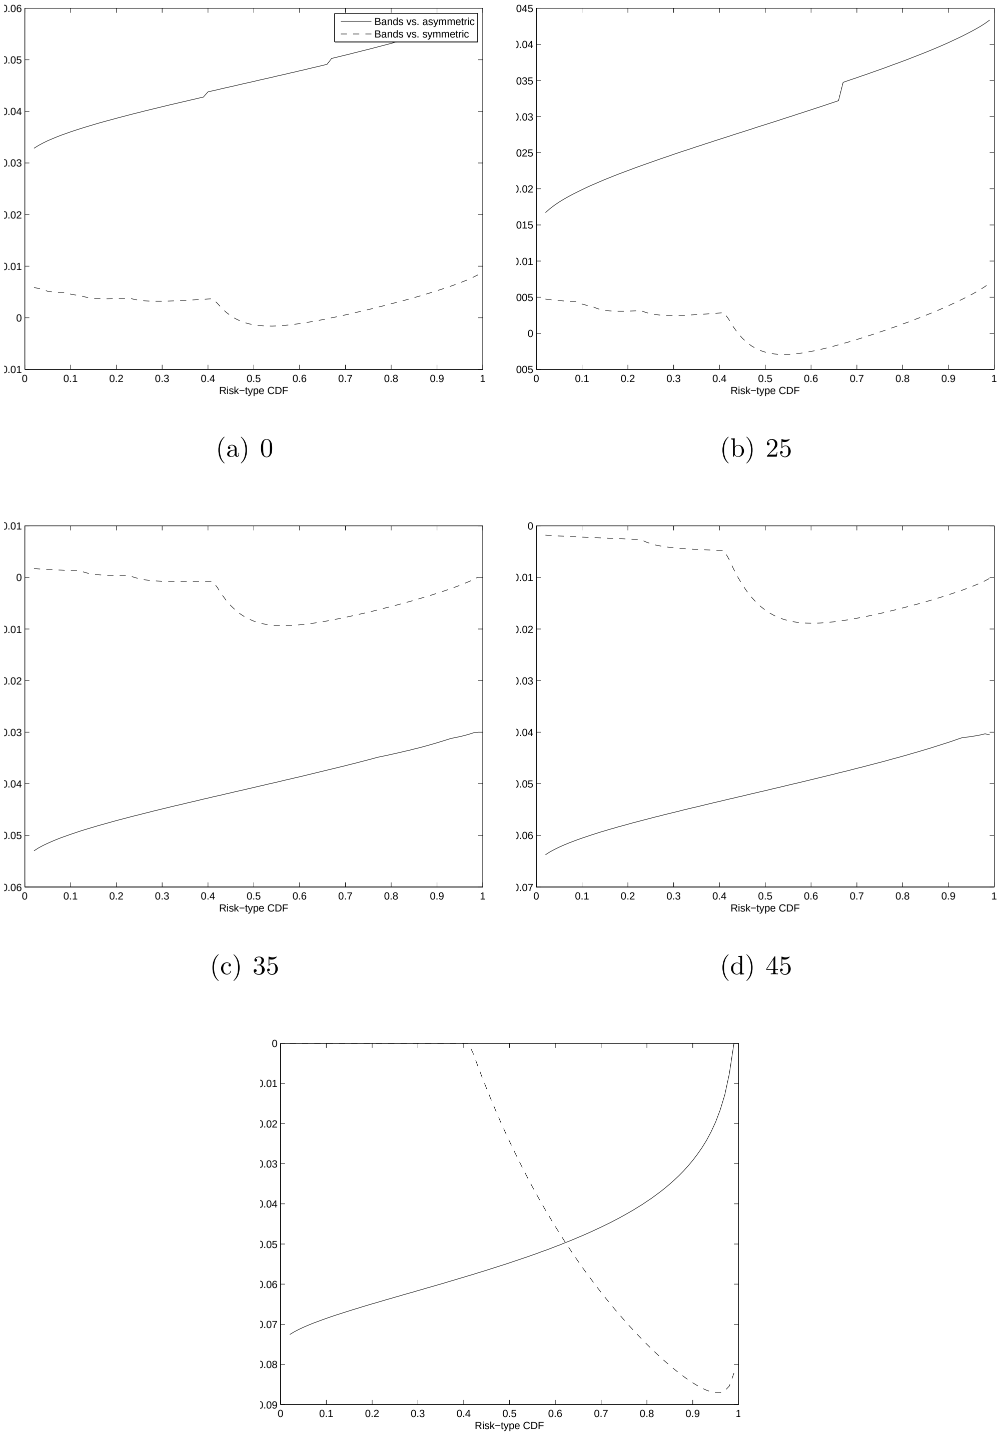

(a) 0 (b) 25

(c) 35 (d) 45

(e) 55

Figure 3: Age-specific certainty equivalents between information regimes with age-group transitions.

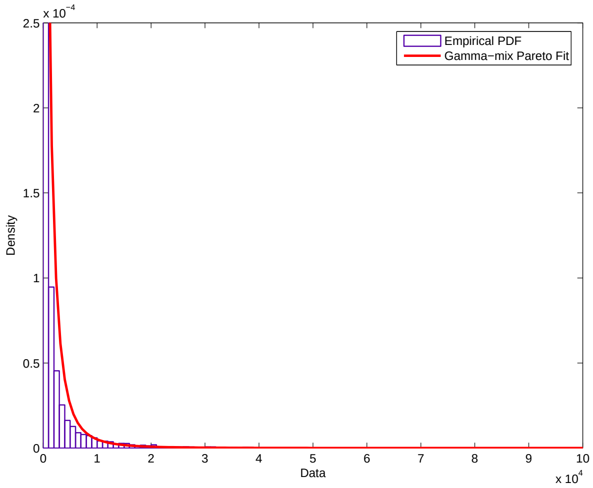

Figure 4: Empirical PDF of Charges and a Gamma-mix Pareto fit

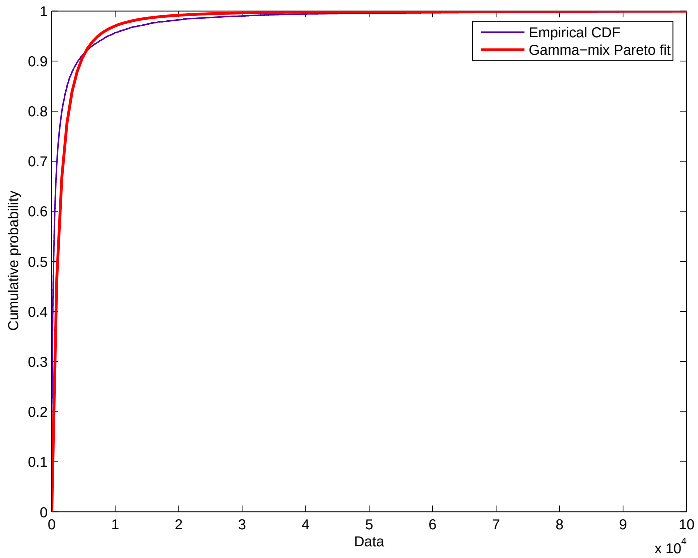

Figure 5: Empirical CDF of charges and a Gamma-mix Pareto fit