Eric Schultz
uniquename: szulc
szulc@umich.edu

Professor Sami
SI 683: Reputation Systems
Final Report
Feb. 28, 2007

A Reputation System for Health Insurance

Introduction

This paper presents a model of a reputation system for solving problems that arise in health insurance. Information asymmetries plague health information and ultimately lead to typical information economic problems, moral hazard, adverse selection, and lack of incentives. We'll only look at a very stylized portion of these issues that occur between the health insurance recipient and the insurance system. This is to say that there are many more stakeholders (e.g. doctors, hospitals, etc...), but that is a far more complicated matter to be tackle in this paper. Some future extensions addressing this are presented in the discussion section of this paper.

The moral hazard that could occur is if insurance recipients are penalized (i.e. higher premium) for being high-risk. We can expect that high-risk individuals will misrepresent their risk status in order to save money. By the time risk status can plausibly be verified it is too late, and costs have been incurred (Campbell 1995). If we are in a situation where costs cannot be covered, premiums have to be raised. If they are raised too far, low-risk individuals will be better off exiting the system, and, thus, creating adverse selection. When low-risk individuals leave, premiums have to be raised again, and more individuals are better off exiting the system. We can expect that this could potentially spiral down until only high-risk individuals are in the system, and a lot of people are left without insurance (Akerlof 2002). If we could somehow sustain the system with a fixed premium, high-risk individuals are not internalizing the externality that they impose on low-risk individuals. High-risk individuals have no incentive to convert to a low-risk status, and we have lost cost-reducing opportunities (Campbell 1995).

Base Model

Let's assume that there are two types of individuals, low-risk and high-risk. Low-risk individuals engage in behaviors that lead to a lower probability of incurring a health cost than high-risk individuals. We'll assume that the health cost that could be incurred is the same for both individuals. In other words, the only difference is the probabilities between individuals. The

individuals are also not selected by nature as one type or the other. This means that at any time an individual can convert from one type to another. For instance, a smoker quits smoking or a non-smoking decides to start.

Let's assume that there is a benevolent monopolist with no marginal administrative costs. The monopolist collects and assigns insurance premiums and pays for health costs. This assumption is subtle. Unlike a government program where everyone has to participate, we want to treat individuals as autonomous and free to choose whether s/he or wants insurance. This ensures that the system is robust against moral hazards and adverse selection while giving incentives to reducing health costs. A government program would merely redistribute wealth from the wealthy to the poor and does not incentivize cost-reducing behavior.While a government program would not have adverse selection, we might find this infringement on individual choice unpalatable.

Let's assume that individuals are risk-averse and all individuals have the same utility function. They prefer to pay more money than the expected cost for not having to bear the risk of incurring a health cost. This is reasonable assumption in that there would be no demand for insurance if everyone was risk-neutral.

Let's also assume that we can verify at the moment a health cost has been incurred whether an individual has misrepresented a characteristic correlated with risk-profile. For instance, a smoker states that s/he is not a smoker but when lung cancer befalls him or her we can verify that s/he is indeed a smoker. We could simply punish this individual at that moment by making him or her pay back the money that was saved through misrepresentation. Unfortunately, we'll have difficulty assessing how long misrepresentation has been occurring. We also might expect that individuals would engage in collusive or exploitative behavior to avoid this harsh penalty. In other words, it would be better if an individual preferred being honest to manipulating the system. We might also imagine that we would like to cover the cost no matter what an individual's choice is (i.e. not punish so severely that it undermines the point of insurance). Lastly, because health insurance spans over such a long time horizon, punishment at the end of the time horizon means that we have lost many cost-saving opportunities gained from converting a high-risk individual to a low-risk individual.

Stylized Model

Let's imagine that everyone's disutility function is $U = f(x^a)$ where $a > 1$ and x is the loss of money. The expected disutility for risk of losing money is $EU = P_r * U + (1-P_r)*0$ where P_r represents the probability that an individual will incur a cost, 0 is the state that an individual does not suffer a health cost (i.e. no disutility). P_r for low-risk individuals is less than P_r for high-risk individuals and will be denoted as P_l and P_h. For instance, $P_l = .01$ and $P_r = .05$ means that a low-risk individual has a 1% chance of incurring a cost and a high-risk individual has a 5% chance. Because everything in the system is a cost we are concerned with disutility which justifies why

a>1. In other words, each additional dollar paid feels worse than last one, and individuals prefer to minimize disutility (Kahnemann and Tversky 1979).

Let's add some numbers here. a=2, the expected health cost is \$10,000, $P_l = .01$ and $P_h = .05$. The expected cost of the individuals are \$100 (low-risk) and \$500 (high-risk). If we had complete information we could simply collect this money and all costs would be covered. Unfortunately, the monopolist is at an information disadvantage and cannot separate the two individuals. This means that high-risk individuals can exploit the system by misrepresentation and pay less, and the monopolist will be paying out more than what is collected. We expect the monopolist would then go bankrupt and leave everyone without any insurance.

Fortunately, individuals are risk-averse and are willing to pay more money to avoid risk than what the expected cost would be. We can use this knowledge to create a sustainable system. From the stylized example, a low-risk individual's $EU_l = .01(10{,}000^2) = 1{,}000{,}000$ and high-risk is $EU_h = .05(10{,}000^2) = 5{,}000{,}000$. The maximum premium (or revenue) we can collect before an individual's disutility is greater than the expected cost without insurance is $R = EU^{1/2}$. $R_l = 1{,}000{,}000^{1/2} = \$1{,}000$ and $R_h = 5{,}000{,}000^{1/2} = \$2{,}236.06$. First, we should notice that we can charge high-risk individuals a higher premium than low-risk individuals. Secondly, under certain risk conditions both premiums are greater than the expected health cost. This means without any information we could charge the lower premium and cover costs without inducing low-risk individuals to exit the system. Secondly, because Revenue>Expected Cost, we can use the extra profits to create incentives. We actually want a monopolist to extract as much profit as possible, and we desire benevolence so that those profits will be redistributed to create incentives.

Mechanisms

The monopolist collects claims, makes payments, and redistributes profits in recurring cycles. Let's assume that this cycle happens monthly. At the beginning of each cycle every participant is endowed with a TrustShare. A TrustShare is similar to an investor purchasing a share in a publicy-traded business. Just as shares increase when a business becomes more profitable, a TrustShare becomes more profitable as the trusted individual increases trustworthiness. An individual (truster) can allocate a TrustShare to any other individual (trustee) that s/he wishes. We'll call this allocated TrustShare a TrustLink from the trustee's perspective. In other words, individuals allocate TrustShares but receive TrustLinks. At any moment a truster can reallocate his or her endowed TrustShares to any other set of trustees. In addition, a truster can allocate multiple or all TrustShares to one trustee.

At the end of each cycle, the monopolist redistributes a portion of profits to each individual. The recipient keeps 25% of the redistributed profit, and the rest is divided evenly between his or her trusters. For instance, A,B,C have allocated 10 TrustShares (2,3,5 respectively) to individual D.

$100 is redistributed to D at the end of a business cycle. D keep $25, $14 is passed to A, $21.5 is passed to B, and $37.5 is passed C. In other words, each truster receives (.75R/t)*n, where R is the money received, t = total shares, and n = number of TrustLinks.

The amount of revenue a trustee receives is proportional to the number of TrustLinks s/he has. For instance, let's imagine 3 individuals with TrustLinks, A:10, B:7, C:3 and $2500 of total profit is to be redistributed. A receives $1,250, B receives $875, C receives $375. In other words, R=(Profit/t)*n, where t = total number of TrustLinks in the system, and n is the number of TrustLinks an individual has. Even though an individual gets a constant percentage from redistributed profits, an individual still prefers more TrustLinks. More TrustLinks means a bigger piece of revenue that 25% is taken from. At this point, though, it's uncertain if an individual would allocate a TrustShare to a more trusted individual. In other words, would it be more profitable to allocate a TrustShare to someone that already has many (small percentage of large revenue) or someone that has none (large percentage of small revenue)? We might consider that this is a nice feature in that individuals can weigh the trade-off between exploitation and exploration.

We need to add a couple more mechanisms in order to make the system work effectively. In it's current state, the system would induce a race to get as many TrustLinks as possible with no correlation to high/low risk profiles (i.e. it's a popularity contest). Remember that we can verify after a claim whether an individual has misrepresented his/her characteristic. If an individual is caught misrepresenting him/herself revenue received from TrustLinks (the 25%) during misrepresentation is taken away and redistributed. In addition, revenue that individuals received from the corresponding TrustShares during misrepresentation are taken away and redistributed to everyone else using the aforementioned allocation function. Individuals do not lose their endowed TrustShares.

Do high-risk individuals want to misrepresent their risk-status? They get the same revenue from TrustLinks no matter whether they are high or low-risk. So, it doesn't seem that high-risk individuals would prefer to take the risk of misrepresentation that has no additional benefit. Will an individual allocate TrustShares inconsistent with his/her belief of a trustee's trustworthyness? Consider an individual's choice to allocate a TrustShare. The individual is presented with four options, truthfully allocating to low-risk, truthfully allocating to high-risk, falsely allocating to low-risk, falsely allocating to high-risk. Assuming all four options have the same number of TrustLinks, then the revenue received via the TrustShare is the same after redistribution. The choice the individual has to make when allocating is to be truthful or not. Considering risk-profile makes no difference. So, the question presented is does one want to be in a risk or non-risk situation even when the amount of revenue is the same? The truster will prefer no-risk because s/he is risk-averse.

The last mechanism is to bundle the health insurance with a good that has a value negatively correlated with the value of insurance. This is to say that low-risk individuals value insurance less but value the bundled good more, and vice-versa (Eppen, Hanson and Martin 1991). We propose this good to be a 401k account. All revenue earned by individuals via TrustShares and TrustLinks is placed in a 401k account and cannot be withdrawn until retirement. It's plausible that low-risk

individuals would place a greater value on long-term investment, hence, why they are low-risk. High-risk individuals value having immediate health care costs covered but are less likely to value long-term investment as much. One nice characteristic that this introduces is the inability for an individual to game the system in order to collect quick cash. Because the risk of losing revenue is present over such a lengthy time frame (i.e. the shadow of the future) we might expect that individuals will cooperate and not manipulate the system (Bó 2005).

Creating this bundle means that we can charge the same premium to all individuals. In fact, we can start to increase premiums to reward trusted individuals. Assuming that the system described thus far is working as expected, we can calculate the appropriate premium. For instance, if the system is all low-risk individuals we charge $1000. If the system is all high-risk individuals we charge $2,236.06. More than likely, the population will be mixed, and we can charge somewhere in between. We have to be wary of adverse selection occurring. Since trusted low-risk individuals are receiving revenue, we can calculate whether a raised premium will be ameliorated by revenue received via TrustShares and TrustLinks. Why might we want to increase premiums? Experiments have shown that punishment can be effective in reducing free-riding in public goods (Ernst and Gachter 2000). We'd like to confront everyone with the fact that high-risk individuals are causing premiums to raise. We can then expect that low-risk individuals might use a tactic of removing TrustShares from high-risk individuals as punishment. This would then create a moderate incentive to high-risk individuals to convert to low-risk. At the very least, it would start to make high-risk individuals pay for the externality that they impose. This does add a little ambiguity in that there may be trustworthy high-risk individuals not recognized as such. This is a fair trade-off. Basically, we only care if a low-risk individual is trustworthy because s/he is the one paying more than expected. High-risk individuals are always paying less than desired.

The Computer-mediated System

Undoubtedly, this model would not be realistic without being mediated by computers. Fortunately, it lends itself nicely to a computer interface. First, we would simply allow individuals to create a profile. For the sake of the model, an individual would only have to indicate whether s/he is a smoker. It's easy to imagine, though, that many other characteristics could be captured in the profile. An individual can add text and photos in order to publicly make his/her case of trustworthiness. An individual would be presented with a visualization of his/her trust network. Lastly, profiles can be searched and history of TrustLinks/TrustShares can be viewed by anyone. This would allow individuals to discover new trustworthy individuals or discover manipulators (e.g. if TrustShares are for someone are dropping dramatically). The system should work as expected with only these simple interface elements.

Conclusion & Discussion

This paper presented a health insurance model that ameliorates moral hazard and adverse selection by allowing individuals to determine the trustworthiness of other members. Incentives for cost-reducing behavior can be achieved by allowing individuals to punish high-risk individuals. All of this was accomplished without having to create different premiums or inhibiting autonomy.

This model only covers a small portion of the issues present in health insurance. For instance, we assumed that individuals had an appropriate budget to pay for insurance. In reality, some people who need insurance the most are poor. Poor individuals might also be statistically high-risk because they live in rough neighborhoods, and we might find it unpalatable to penalize an individual for risk that is attributed by nature. A more robust trust model might allow for variations in income and high-risk caused by nature not choices. A more disturbing issue is that we assumed that high/low-risk could be correlated with an observable characteristic (e.g. smoking), but in reality things are not as straight-forward. For instance, three people contract HIV. One is sexually conservative, one is sexually liberal (i.e. many partners), and one is a sex worker. We might have a hard verifying whether an individual misrepresented him/herself because HIV only somewhat correlates with sexual activity. In other words, we can't infer that someone that contracts HIV was or was not lying. This paper also assumed that health recipients are the only stakeholders. Hospitals and doctors represent more moral hazard issues. Is there incentives for innovation and non-excessive treatments? It's plausible that the model could be extended to address a bigger portion of the issues that plague the health system.

References

Akerlof, G. (2002). "The Market for 'Lemons": Quality, Uncertainty and the Market Mechanism". *Economic theory and the welfare state.* 1:308-20.

Bó, P. (2005). "Cooperation under the Shadow of the Future: Experimental Evidence from Infinitely Repeated Games". The American Economic Review. 95(5): 1591-1601.

Campbell, D (1995). "Incentives: Motivation and the Economics of Information". *Cambridge University Press*.

Eppen, Gary D., Ward A. Hanson, and R. Kipp Martin (1991). "Bundling - new products, new markets, low risk." *Sloan Management Review* 32.n4: 7(8).

Ernst Fehr, Simon Gachter (2000). "Cooperation and Punishment in Public Goods Experiments." *The American Economic Review*, 90(4): 980-994.

Kahnemann, D., Tversky, A (1979). "Prospect Theory: An Analysis of Decision Under Risk." *Econometrica* 47(2): 263-291.