Health Economics in the Small, Medium, and Large

A Lecture at the European Health Economics Workshop, Bergen, Norway, May 2008

Ching-to Albert Ma

Department of Economics
Boston University and University of Oslo

email: ma@bu.edu

Acknowledgement: My thanks to Randy Ellis, Jacob Glazer, Ting Liu, Michael Luca, Tom McGuire, Monic Sun, and Reema Vaithianathan for their comments.

1 Introduction

It is my honor to address you here in Bergen on Health Economics. I feel that I am very unqualified to give a keynote lecture on anything, especially on health economics. I am speaking to very esteemed and senior experts in the field, indeed. I felt so nervous that I almost asked my half-brother Yo-yo Ma for help; Yo-yo is certainly very qualified to play a note and a key, and of course to lecture about his trade. The same cannot be said about Albert.

By the way, I should add that Chinese take half-brothers very seriously. Basically, Chinese people who have the same last name are all half-brothers or half-sisters, so I estimate that I have more half-brothers than the number of Scandinavians.

Many years ago, when I started learning about health economics, particularly from Randy Ellis and Tom McGuire, I quickly realized that most health economic problems boil down to either moral hazard, adverse selection, or both. Moral hazard and adverse selection simply will remain as the main ways to think about health economics. Nevertheless, my own thinking in recent years has led me to believe that perhaps the classical moral-hazard-adverse-selection framework should be expanded. I would like to share with you some of these ideas.

You may be amused by the title, and I think that I am, too. What do Small, Medium, and Large mean? They may refer to sizes of cakes, for example:

  

Here, they actually refer to perspectives on health economic issues. Perspective is a subjective word, and of course I will be telling you some personal views. But for what it is worth, the title says something like this. Small refers to individual health insurance, Medium refers to provider-patient interactions, and Large refers to health markets. Maybe you will agree with me that in fact the Small, Medium, and Large classification is not too bad. As you will see, as the perspective becomes bigger and bigger, the departure from the classical moral hazard and adverse selection framework becomes more and more obvious.

2 Health Economics in the Small

Let me begin by the alarming macro fact that the US is spending about 16% of its GDP on health care. It has never been higher, and there is no sign that the growth will slow down. But this observation reminds me of an occasion when a student showed me a data set of a developing country where an average household was reported as spending about 80% of its income on health care. How could an average household spend 80% of its income on health care? There must be something wrong. Either the data were inaccurate, or the definition of health care was too broad. I have a feeling that these days perhaps the term health good means a lot more stuff than it used to.

Are acupuncture, massage therapy, spa treatment, anger therapy, full-body scan, gene test, nose job, body-part augmentation, meditation class, dance class, laser eye surgery, infertility treatment, gastric bypass, Prozac and Ambien health goods? Consumers want some of these goods whether or not they are “sick” or “healthy.” I can be pretty sure that those who buy these goods, with or without insurance, perceive them to be of some good to them. Insurers, however, treat the modern plethora of health goods either with confusion or distaste. There are, therefore, all kinds of disputes, both in North America, and in European countries, about what are and what are not covered by health insurance.

I want to propose a theory that perhaps sheds a little bit of light on this confusion. The idea is that the uncertainty nowadays is not just about health status, whether one is sick or not, or how sick one is, but rather about preferences. I want a pair of glasses over my nose because it gives me the much needed "intellectual" look to give my keynote lecture; a more gregarious professor would think that the same pair of eyeglasses would deter the opposite sex, and therefore opt for laser eye surgery. But when I was a kid, I could not tell whether I would be a self-deprecating professor or a sociable one.

And here is the surprise. I am going to show you that we are all risk-loving when it comes to preferences. If I am offered a fifty-fifty bet about being a self-deprecating professor and a gregarious professor, versus a "normal" professor like Kurt, who, as we all know, is just the right mix of the two extremes, I will always go for the lottery. And, "I" am not alone. I am going to show you that you will do the same, even if your utility function is concave in all the consumption goods!

Here is a parsimonious model. Let there be N goods, and I write a consumer's utility function as $\sum_{i=1}^{N} \theta_i \; U(q_i)$, where U is an increasing and concave function, and $\theta = (\theta_i)_{i=1,\ldots,N}$ a vector of weights. Suppose that the prices of these goods are p_i, $i = 1, ..., N$, and that the consumer has income M to spend on these goods. You are all familiar with the standard consumer problem: choose q_i, $i = 1, ..., N$ to maximize the utility subject to the budget constraint.

I will not bother to write down anything but the indirect utility function:

$$V(\theta; p, M) = \max_q \sum_{i=1}^{N} \theta_i \; U(q_i) \quad \text{ s.t. } \quad \sum_{i=1}^{N} p_i q_i = M.$$

If you take out your first-year graduate micro textbook, you will find out all there is to know about the indirect utility function. It is continuous, decreasing in prices, increasing in income, and quasi-convex in prices and income, etc.

What the textbook does not tell you is that the indirect utility function is convex in θ, the preference weight parameters! I will show an example by the familiar Cobb-Douglas utility function: $U(x_1, x_2) = \theta \ln x_1 + (1 - \theta) \ln x_2$. Normalize income M at 1 , and let the prices be p_1 and p_2. Now consider three different values of θ: 0, 0.5 and 1. Here are the optimal bundles and indirect utilities:

	x_1	x_2	V
$\theta = 0$	0	$\dfrac{1}{p_2}$	$-\ln p_2$
$\theta = 0.5$	$\dfrac{0.5}{p_1}$	$\dfrac{0.5}{p_2}$	$\ln 0.5 - 0.5 \ln p_1 - 0.5 \ln p_2$
$\theta = 1$	$\dfrac{1}{p_1}$	0	$-\ln p_2$

Because $\ln 0.5 < 0$, the average of the indirect utilites at $\theta = 0$ and $\theta = 1$ is larger than the indirect utility at $\theta = 0.5$. The indirect utility function for Cobb-Douglas preferences is convex in θ. Indeed, the general proof is elementary, and not worth the back of an envelope.[1]

When V is convex in θ, it follows that

$$\int V(\theta; p, M)\ \mathrm{d}F(\theta) \geq V(\overline{\theta}; p, M),$$

where F is the distribution function of θ, and $\overline{\theta}$ is the mean of θ. When I was young, if I was offered an "insurance" policy to get rid of a 50-50 lottery of becoming the self-deprecating professor and the gregarious professor in order to become a professor like Kurt for sure, I would have refused the insurance. Playing the lottery of life is so much more fun, and you would have to bribe me to become a professor like Kurt:

$$\int V(\theta; p, M)\ \mathrm{d}F(\theta) = V(\overline{\theta}; p, M + \epsilon) \ \text{ where } \epsilon \geq 0.$$

(And I guess now you know why truly rational intellectuals are stupid, smart, crazy, depressed, manic, or all of the above, but never normal!)

---

[1] Let $q(\theta)$ be a solution for any given θ. Let $\theta = \alpha\theta^1 + (1-\alpha)\theta^2$, where $0 < \alpha < 1$. The quantity vectors $q(\theta)$, $q(\theta^1)$, $q(\theta^2)$ all satisfy the budget constraint. By definition

$$\begin{aligned} V(\theta) &= \sum_{i=1}^{N} \left[\alpha\theta_i^1 + (1-\alpha)\theta_i^2\right] U(q_i(\theta)) \\ &= \alpha \sum_{i=1}^{N} \theta_i^1 U(q_i(\theta)) + (1-\alpha) \sum_{i=1}^{N} \theta_i^2 U(q_i(\theta)) \\ &\leq \alpha V(\theta^1) + (1-\alpha) V(\theta^2), \end{aligned}$$

and hence V is convex in θ.

Of course there is no insurance policy that can possibly guarantee preferences, but that is not the point. The point is that we have no tendency to avoid extreme preferences, and that seems to go against the insurance for health risks. The usual moral hazard story is to be revised in light of this tendency. But how?

Let me consider how health insurance will fit into this model. We know that the indirect utility function is quasi-convex in income M. But the indirect utility function is increasing in income, so quasi-convexity does not say very much. Perhaps the textbooks have forgotten to mention that in fact the indirect utility function is concave in income! Again the proof is elementary.[2] When V is concave in M, it follows that

$$\int V(\theta; p, M) \ \mathrm{d}G(M) \leq V(\theta; p, \overline{M}),$$

where G is the distribution function of M, and $\overline{M}$ is the mean of M.

Allow me to define health risks as income uncertainty. Indeed, the consumer is risk averse with respect to income fluctuations. We do have the demand for insurance, insurance against the variation of marginal utilities of income. I have already shown that the consumer is *not* risk averse with respect to preference fluctuations. What does insurance mean when a consumer is risk averse in one sort of variation (such as incomes) but risk loving in another sort of variation (such as preferences)?

One can perhaps speculate that even in this model an insurance policy always takes the form it does: one pays a premium, and the insurance policy pays a subsidy. The consumer does not bear the full price ex post. What has been puzzling me for some time is why some "health" goods are completely excluded. The usual answer is "moral hazard." But doesn't severe moral hazard imply a very small subsidy rate, not

[2] Let $q(M) = (q_1(M), ...q_N(M))$ be the solution for any given income M. Let $M = \alpha M_1 + (1-\alpha)M_2$, where $0 < \alpha < 1$. Consider the bundle $q \equiv \alpha q(M_1) + (1-\alpha)q(M_2)$. Clearly, $p \cdot q = \alpha p \cdot q(M_1) + (1-\alpha)p \cdot q(M_2) = M$, so bundle q satisfies the budget constraint at income M. We have

$$\begin{aligned} V(M) &\geq \sum_{i=1}^{N} \theta_i U(q_i) \\ &= \sum_{i=1}^{N} \theta_i U(\alpha q_i(M_1) + (1-\alpha)q_i(M_2)) \\ &\geq \sum_{i=1}^{N} \theta_i [\alpha U(q_i(M_1) + (1-\alpha)U(q_i(M_2)] \\ &= \alpha V(M_1) + (1-\alpha)V(M_2) \end{aligned}$$

and hence V is concave in M.

outright exclusion? The model of preference *and* income fluctuation perhaps offers one way to explain this exclusion.

The preference parameters ultimately interact with prices. If you have the bad habit of writing down the first-order condition every time a constrained optimization is presented to you, then you will see

$$\theta_i U'(q_i) = \lambda p_i \quad \text{or equivalently} \quad U'(q_i) = \lambda \frac{p_i}{\theta_i},$$

where λ is the multiplier of the budget constraint, or the marginal utility of income. Prices eventually matter because they will have to interact with the preference parameters θ_i. Who knows, if the consumer loves variations in θ, may an insurance policy accomodate that by abandoning subsidies, or even trying to tax consumers? An exclusion is simply the boundary condition that the insurance company cannot impose a tax on a consumer because there are no taxes in the open market.

Maybe this two-dimensional structure allows us to explore some issues in the insurance market that are not quite easy to explain in the usual way. These are all speculations. But I have only been asked to talk about keys and notes, not songs, concertos, or symphonies. So, there you go, here are the keys and notes for a more general model of insurance, where health risks and preference lotteries may matter more than we have so far appreciated.

3 Health Economics in the Medium

It is time to move on to the next size. Physician agency has always been a popular topic for research, in the same way that doctor stories have always been popular in TV shows and movies. The definitive survey is the Handbook chapter by Tom McGuire. There are many variations in the theme. I can mention imperfect agency, collusion, altruism, induced demand, etc. There are some common features, though. First, the provider-patient relationship is often modelled as a coalition, which resolves some conflicts. Second, either the patient, the provider, or both will do something, or possess some information.

On a side note, however, I'll say that times have really changed. Economists have finally opened themselves up to psychologists and sociologists, even neuroscientists and psychiatrists. Ten years ago, if I wrote a paper and let some economic agents be honest, people would laugh....how silly an assumption! Now, recently Tor, Geir and I wrote a paper that showed that there were honest and sincere Norwegian doctors, people would say isn't that obvious, surely Norwegian doctors must be honest and sincere? Suddenly, selfish behaviors have diminished in stature, and economists have grown to respect many kinds of motives. Now you see that every paper on physician agency has an altruistic doctor, not unlike the one Randy Ellis and Tom McGuire postulated more than twenty years ago. But it has taken us that long to catch up with Randy and Tom's 20-10 vision.

Having said the caveat, I would like to share with you some idea about a somewhat different approach: the placebo effect in provider-patient interactions. Many books and articles have been written on the placebo effect.[3] Being seen by someone wearing a white coat with a stethoscope hanging on the neck is already remedial. Being given a shot of anything is already reassuring. Sitting down on a recliner and chatting about whatever already is uplifting.

For my purpose, I simply identify the placebo effect by the postulate that an encounter between a provider and a patient may be beneficial to the patient. The doctor and the patient do not have to do too much of anything. There is no need for altruism in an explicit sense, or even professional and ethical concern. Nor do they have to possess any private information. But the placebo effect is not inconsistent with profit or utility maximization, altruism, ethics, or professional standards, either.

How do I model the placebo effect in patient-provider interactions? Let me start with the formula:

$$\text{Outcome} = \text{Natural course of Illness} + \text{Placebo effect} + \text{Treatment effect}$$

I think that we all understand the natural course of illness (and for all us, eventually it means demise), and treatment. First, I am going to let the placebo effect be random. Seeing a doctor may make a patient *feel* better, if not better. Perhaps, there is a lot of truth in the following quote from Shakespeare's Hamlet

There is nothing either good or bad, but thinking makes it so.

I think that you should see my dentist and psychologist; I am absolutely sure that they are the loveliest doctors for you. But I would never recommend to you the awful meditation teacher I had last year. Of

[3] A very good one is "The Placebo Effect and Health," by W. Grant Thompson, 2005, Prometheus Books, New York. A more "medical" account of the placebo effect is in "Placebo: Mind over Matter in Modern Medicine," by Dylan Evans, 2004, Oxford University Press, Oxford.

course, it's only my thinking that I am telling you. You may think and feel differently. This is one reason why it is thought that the placebo effect is random and only works in about 30% to 40% of the time.

Second, I am going to let the placebo effect be persistent. If you happen to like my dentist and psychologist after one or two encounters, maybe you will continue to like them. Sometimes, you may even train yourself to have persistent thoughts about a healing relationship. But I am going to assume that once the placebo effect is realized, it remains unchanged.

With these two components, I am now going to describe to you a 2-by-2-by-2 model: two patients, two doctors, and two periods, and with these three 2's, I add one distribution function, the placebo effect. Then I put in matching, learning, switching, and pricing, all the usual stuff. Insurance, and managed care will come in just a minute.

I do not have to identify the two patients, so no names for them. Each patient wants one unit of treatment from a doctor in each of two periods. The two doctors are called $D1$ and $D2$. The unit cost of treatment is a constant $c \geq 0$, for each doctor. A patient must see a doctor. An initial encounter may happen in period 1: a matching occurs, and a doctor-patient matched pair confers a benefit of α to the patient, and the patient and the doctor learn about this. This is the placebo effect. This benefit is distributed according to F on $[\underline{\alpha}, \overline{\alpha}]$, with density f. Patients and doctors are all ex ante identical.

After the initial encounter, in period 2, the patient may stay with the doctor who has treated her in period 1. In this case, the benefit α that has been drawn in period 1 is again the benefit the patient receives. But, the patient may switch to a new doctor, and in this case, a new value of α will be drawn from F. I have described matching, learning and switching already.

Look at this patient and the doctor. See how happy they are. I bet that they have drawn a very high α

We can be sure that she will find him for the second surgery. But if your doctor wears this button

I suppose that your α is low and you will be prepared to look for another.

What about pricing then? I use a long-term pricing structure. A doctor sets a price before any patient comes along, and sticks to it. Let p_i be the price set by D_i, $i = 1, 2$. If a doctor sees a patient in a period, he gets $(p_i - c)$, and if he sees two patients a period, he gets $2(p_i - c)$, and so on. If a patient sees a doctor and obtains benefit α, her payoff is $\alpha - p_i$. Patients are risk neutral, so if she does not know α when she sees a doctor, her expected payoff is $\mathrm{E}(\alpha) - p_i$.

I do not allow quantity discount, so that prices stay the same no matter whether the patient receives one unit or two units of treatment from a doctor. In other words, I do not allow prices to vary over time, so that a doctor's price remains the same over time. Old patients and new patients pay the same price. Complex pricing strategies seem harder to deal with.

The extensive form is as follows:

Stage 0: Each doctor chooses a price. Let Doctors $D1$ and $D2$ choose prices p_1 and p_2, respectively.

Stage 1: Each patient picks a doctor for a first visit. In this first visit, a value of α is drawn. This value of α will be the patient's benefit if she continues with the doctor. The patient pays the doctor p_i.

Stage 2: Each patient may continue with the doctor she has selected in Stage 1, or switch to another doctor. If she switches to a new doctor, she pays the new doctor p_j, and then a new value of α is drawn.

This game incorporates quite a few things. First, as I have mentioned, there is the placebo effect. Doctors do not really do anything other than set prices at the beginning of the game. Second, the placebo effect is random and persistent. After a first visit, a patient pays the price and learns the specific benefit from a doctor. The benefit remains the same if she continues with the doctor for a second visit. Third, switching is possible if the patient wants to try another doctor, but again she pays the price for the next first to see if that will be a good match.

Ex ante doctors are identical. After an experimentation, the match outcome is learned by the patient. Doctors are no long identical. If the match does not quite work out, the patient may try another match. In other words, "products" are ex ante identical, but ex post differentiated. The model is ex ante Bertrand, ex post Hotelling.

A consumer's decision in stage 2 is based on the comparison of her current benefit α at price p_i with Di and the expected benefit $\mathrm{E}(\alpha)$ with Dj at price p_j. She switches if

$$\alpha - p_i < \mathrm{E}(\alpha) - p_j.$$

A consumer also anticipates her best choices at the beginning of the game when she picks between doctors for the first match. The expected utility from choosing Di is

$$\mathrm{E}(\alpha) - p_i + \int_{\alpha} \max\left[\alpha - p_i, \mathrm{E}(\alpha) - p_j\right].$$

I can now make the first observation: in stage 1, a consumer always picks the doctor with a lower price. Ex ante doctors are identical, so a consumer optimally picks Di if and only if $p_i < p_j$. A doctor who prices high loses all consumers in stage 1. But, the high-price doctor does not make zero profit! Although no patient comes to him in stage 1, some may be so disappointed with the low-price doctor in stage 1 and come to him in stage 2. Ex post a doctor has a downward sloping demand function. In stage 2, he attracts more patients by lowering the price in a continuous fashion.

Lemma 1 *There is no pure strategy equilibrium.*

Such a simple game, and you need to solve for complicated formulas for the mixed strategy equilibria! Maybe you have already figured out why pure strategy equilibria fail to exist. First, no doctor will charge at marginal cost c. Such a low price will never make any profit. A doctor may simply charge a higher price, let the rival doctor take all consumers in period 1, and then, so to speak, make a killing in period 2, capturing patients who become disappointed. Second, there are no equilibria where both doctors charge an identical price above marginal cost. This is due to the standard Bertrand argument. By slightly undercutting the rival, a doctor makes substantially more profits in the first period to compensate for the very small loss in the second period.

Perhaps more surprising is the nonexistence of asymmetric pure strategy equilibria. If you postulate that $D1$ charges p_1 and $D2$ charges p_2 with $c < p_1 < p_2$, you will find that their first-order conditions are inconsistent with each other. It's a contradiction.

What does this have to do with health insurance and manage care? Recall my formula:

$$\text{Outcome} = \text{Natural course of Illness} + \text{Placebo effect} + \text{Treatment effect}$$

I have modelled the placebo effect. The treatment effect can be handled by the usual "action" story. Let Di spend some costly efforts on the patients. The insurer-payer now has an incentive problem to solve: how to motivate doctors to work harder when they compete against each other in the extensive form above.

Doctors already may have some incentives to work. To compensate for a bad match, a doctor can work harder. If I don't like my psychologist or my psychologist does not like me, we can work harder to learn to like each other. Maybe the psychologist can dig deeper for the etiology, or find out which books will be helpful to me. Maybe a dentist can achieve better outcomes by more efforts.

Health insurance means consumer copayment less than 100%. Managed care may be modelled by network. Suppose that $D1$ is a doctor within a managed care network, but $D2$ is not. Consumers may face different copayments when visiting $D1$ and $D2$, and $D2$ may be out of the network entirely. For example, $D1$ is a conventional, family doctor, while $D2$ is an alternative medicine provider. With all the symmetry assumptions on cost and benefits, then all patients will first visit the conventional doctor, but if it does not work out, some patients will try alternative medicine, which is not covered by insurance.

We can also put insurance and effort incentives together. We can also put managed care network and incentives together. In other words, the model without any price subsidy or preferential treatment on providers can be adapted for these considerations.

Again, I am afraid that I have some notes and keys, not a song. I have told you that in the model without any interference from insurers, there are no pure strategy equilibria. The mixed strategy equilibria have yet to be constructed. Some asymmetry would be of interest. We do expect conventional doctors to incur a higher marginal cost than alternative medicine providers, so it is not without interest to let these costs be different. Overall, the model is different from the conventional Bertrand and Hotelling models, and the result already shows. What else is there? I hope that with some time and effort, I can tell you more later.

4 Health Economics in the Large

In every country worth its name, the government plays a significant role in the health market. A completely private health care sector doesn't seem to exist. Therefore, health economics in the Large must deal with interactions between private and public sectors. As a result, policy talks are unavoidable. But these are turbulent times for policy discussions among economists—although the public and politicians have not had similar difficulties. The economists' difficulty comes from behavioral economics, which, through importing some ideas from psychology, biology, and medicine, has questioned the real meaning of welfare. Indeed, how can economists enlighten policy makers when economists themselves cannot agree on the meanings of such words as "better off" and "worse off?"

I cannot address you on behavioral health economics. I myself find it more rewarding to read psychology rather than the economist's version of psychology. I am going to stick to something more friendly, for example, prices and quality, in the public sector and in the private sector. But prices are not enough. Everywhere we turn, we find that in the health sector, there are two kinds of policies: quantity and price policies. Rationing and subsidies are here to stay in the portfolio of public policies for the health sector.

Rationing policies may take many forms: waiting time, waiting list, free supply for poor consumers, priority according to illness severity, etc. I will not go into details here. Generally, they affect quantities. Subsidies can also take many forms: user fees that are based on income, quality charges, and cost sharing. Generally, they work through consumers' willingness to pay.

I will sketch out a big model, since we are dealing with the Large. It consists of a three-dimensional description of consumers. Let w denote consumer's wealth or income, v consumer's valuation of quality, and γ the cost parameter of provision. In this model, each of these dimensions follows some distribution: consumers have different incomes, various valuations of quality, and may require different service intensity for health care.

I will let a consumer have at most one unit of health care, but this is only a matter of semantics, because I allow the quality of health care to vary. Alternatively, a consumer may have many units of health care, each with a fixed quality, but for presentation I will use the unit demand language. I write a consumer's utility as

$$U(w) + vq,$$

where q is the quality of a unit of health care. The cost is of no direct concern to the consumer, but of course the public and private sectors have their keen eyes on cost.

Since the advent of the modern regulation economics, it is customary to apply the mechanism design approach, sometimes almost mechanically. I do not subscribe to that now. Most regulatory problems are quite complex, and the direct revelation game may just miss the point. So no revelation principle today. Further, I will not even bother to write down welfare indexes, so not much optimization to do either. Instead, I will just set up some policies.

I have defined three variables: w, v, and γ. There are only two sectors. So this is going to be a sequence of 2-by-3 models. But what are the questions that can be explored? Mostly, I am going to talk about quality competition, but many other issues can be discussed in this model. They include crowding out, selection, and generally interactions between the public and private sectors.

You will have noticed by now that the quality I am talking about is of the vertical differentiation sort. As much as I admire Hotelling, today I will not use a horizontal differentiation model for the health care sector. I mostly think that a consumer wants to be treated by a good doctor, in a good facility, be it public or private. Locations and office hours are the most natural candidates for horizontal differentiation, but I will ignore them here.

Let me specify the cost by $\gamma C(q)$, so γ is a cost parameter. Supplying a unit of health care at quality q to a consumer with cost parameter γ requires a cost of $\gamma C(q)$. This is the component for selection concern. All patients are costly, but some are more costly than others.

A second concern is a consumer's willingness to pay. In this setup, this can be a complicated matter. How much is a consumer willing to pay for a unit of health care delivered to him at quality q? I can find this by solving for p in the following equation:

$$U(w - p) + vq = U(w).$$

You can easily verify that p is increasing in w, v, and q. A wealthy patient is willing to pay more, a quality-minded consumer is willing to pay more, and health care of a higher quality commands a premium.

Let me begin with the simplest case. I am going to let U be linear. The wealth variable becomes irrelevant. Now we are left with only two: v and γ. There are two sectors, so it is "natural" to consider asymmetric information of the following sort: let the public sector observe v and let the private sector observe γ. Each sector holds on to this private information.

Knowing γ, the private firm considers choosing a price-quality pair to offer to consumers. Let $(p(\gamma), q(\gamma))$ be the price-quality pair when the firm observes γ. For the time being, suppose that the public sector is inactive. Then a consumer buys if and only if

$$vq - p > 0 \quad \text{or} \quad v > \frac{p}{q}$$

which happens with probability $1 - F(\frac{p}{q})$ with F being the distribution function of v. The corresponding profit is

$$\left[1 - F(\frac{p}{q})\right][p - \gamma C(q)].$$

Again, if you cannot resist the temptation of choosing p and q to maximize the profit, and writing down the first-order conditions, you get:

$$\frac{p - \gamma C(q)}{p} = \frac{1 - F(\frac{p}{q})}{f(\frac{p}{q})\frac{p}{q}}$$

$$\frac{p}{q} = \gamma C'(q).$$

These are the familiar marginal-this-equals-marginal-that conditions. The first one is a relationship between the price-cost marginal and the elasticity of demand. The second one is the profit-maximizing choice of quality. For the marginal consumer, the valuation of quality equals to the marginal cost of quality:

$$v = \gamma C'(q).$$

The firm's optimal choices are all functions of the cost parameter γ. Keep this in mind. The public sector observes the cost and offers a price-quality pair accordingly. You can perform some comparative statics of the profit-maximizing price and quality with respect to the cost parameter, but I won't bother here.

I have assumed that the public sector is inactive, but the same sort of optimization will be carried out when the public sector is active. But what is the public sector going to do? I am not going to let the public sector observe the cost parameter γ. Instead, I let the information about valuation v be available to the public sector, but only v, no γ. This is the my line of argument. The public sector observes some information, while the private sector observes some other information. The two sectors then interact.

The public sector can use the valuation information to do a couple of things. But first let me ask why the public sector should do anything at all. The reason is that the prices and qualities in the private market may not be efficient; they may not be what the regulator wants. Clearly the price is too high, higher than marginal cost. Then the quality is chosen to maximize the net benefit of the marginal consumer, and again this may not be socially desirable. Quality is like a public good, and it is the average valuation that matters, not the valuation of the marginal consumer.

What can the public sector do? Let me begin with a subsidy policy based on v. For a consumer with valuation v, the public sector can make available the good at quality q for a price s. Generally, it can make the subsidy and quality functions of valuation v, so I write $q(v)$ and $s(v)$. Now suppose that the consumer with valuation v buys from the public sector. The total consumer surplus due to the subsidy is $vq(v) - s(v)$, and the expected public expenditure is $\mathrm{E}(\gamma)C(q(v)) - s(v)$.

I have not let the public sector make use of the cost information; I have not let the private sector make use of the valuation information. One piece of information for each sector. Now I can put them together. Given the subsidy-quality policy in the public sector and the price-quality schedule in the private sector, a consumer now can pick between the suppliers. Consumer (v, γ) chooses the public supply if and only if

$$vq^{\text{public}}(v) - s(v) > vq^{\text{private}}(\gamma) - p(\gamma).$$

This determines the distribution of consumers across the two sectors. The following diagram illustrates one such distribution:

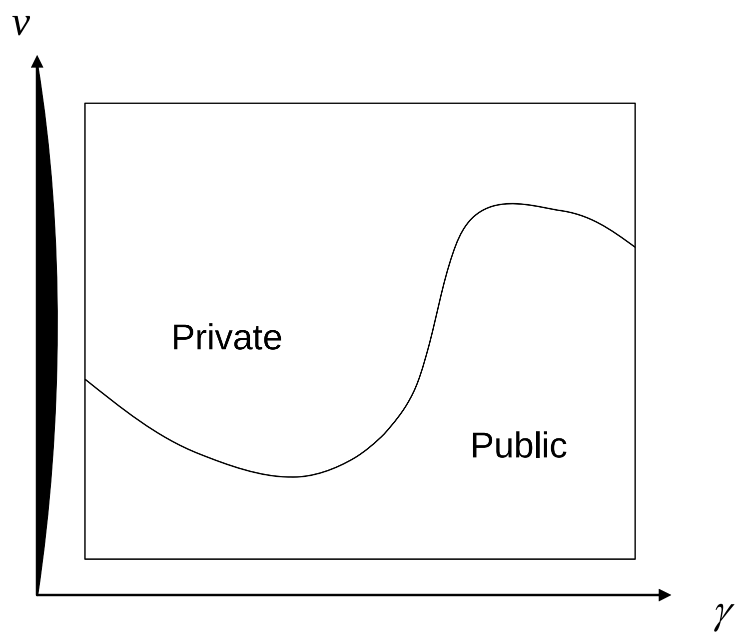

where the line is given by the equation $vq^{\text{public}}(v) - s(v) = vq^{\text{private}}(\gamma) - p(\gamma)$.

I let the public and private sectors "compete" for patients. But each sector's instrument depends on its own private information. Here is a model where players use different kinds of information to compete against each other. The public sector probably does not maximize profit, but it sure wants to make use of the cost information, and it can only do it through the private sector, via its price and quality responses.

I can alternatively let the public sector pick a rationing-quality policy. Suppose that the policy is written in terms of a provision of a unit of the good at quality q to consumers with valuation v at zero cost. For example, rich consumers get rationed out of the public sector, and must purchase from the private sector if they are willing to, while poor consumers get the goods for free. Again, there will be a price-quality reaction by the private sector.

Under either subsidy policy or rationing policy, it is a model of interaction between a public sector and an imperfectly competitive private sector. Each sector has its own private information, but not enough to achieve its grand objective. Consumers are reacting to all available options. The health care setting is appropriate since one usually thinks that health care decisions are multi-faceted.

Within this model, up to now I have considered the case of a linear U function. But I can alternatively consider a case of a concave U function. To keep things simple, one may set v at 1. Here, the utility of the consumer becomes $U(w-p)+q$. Then the willingness to pay will be given by the equation $U(w-p)+q = U(w)$. I can even contemplate interchanging the information structure, so that the public sector may observe cost but the private sector may observe wealth.

Much thought and work need to be put into the above structure to specify the game completely. Again, I only have some notes and keys.

5 Conclusion

It is my privilege to have an opportunity to deliver a lecture on health economics, a subject that I have enjoyed thinking about. I have attempted a somewhat unconventional delivery. I have not given you a survey of some ideas; I thought that life is too short for surveys. I have not given you theorems or propositions; I needed more time for them. I have not attempted to present a research paper either.

Instead, I have explored with you several ideas. As you have seen, they are speculative ideas that have not been worked out very much. But I hope that you will agree with me that these ideas have deviated from the conventional moral hazard and adverse selection framework.

I have not come across papers that deal with the Small, Medium and Large issues in the particular ways that I have formulated them. Perhaps, the reason that no such literature exists is because these ideas are uninteresting. But maybe, just maybe, I can imagine that some new results can be developed, along these lines, but that's for the future. As I said, I have some keys and some notes, no sonata, concerto, or symphony. I have very much enjoyed preparing this lecture, and to me that is what really counts.