Valuing life

Hiroaki Sakamoto

January 9, 2016

1 Hedonic price model

In the classical theory of consumers, individual's decision involves a choice of *quantity*. As we learned in the introductory microeconomics course, consumers decide how many goods, say apples, they buy in the market. By doing so, we implicitly assume that apples are a homogeneous good, namely, every apple is of the same quality. Although this is a good description of reality, people's decision is in some cases better characterized as a choice of *quality*. In the housing market, for instance, you can find different types of houses with varying qualities. Housing is a typical example of what we call a *differentiated good* — a category of good in which each product has different quality or characteristics. Assuming that you buy only one house at a time, then the question is not how many, but which type of house you should choose. In the competitive market, a house of higher quality will be sold at a higher price. Hence, there exists a natural trade-off between the price you pay and the quality of good you obtain.

In order to motivate the analysis that follows, suppose that we observed a stable relationship between the housing price and the ambient air quality. For example, let's say that the price and the ambient air quality are positively correlated as in Figure 1. For simplicity, we assume that other relevant characteristics of housing are fully controlled for. In other words, all the houses included in our data set are identical apart from their ambient air quality. Using this data, we can estimate the *hedonic price function* or *hedonic price schedule* — a mapping that relates a housing price to a particular level of ambient air quality. This function, if properly estimated, reveals consumer's preference about air quality. In particular, the price function must contain information about how much people are willing to pay for improving the air quality around their houses.

1.1 Consumer's decision

To formalize the idea, denote by $q \in \mathbb{R}$ the quality of a differentiated good. The price of the good depends on its quality. So let $p(q)$ be the price function. We note that $p(q)$ is in general non-linear in q. To simplify the analysis, we assume that there exists a sufficiently wide variety of products traded in the market

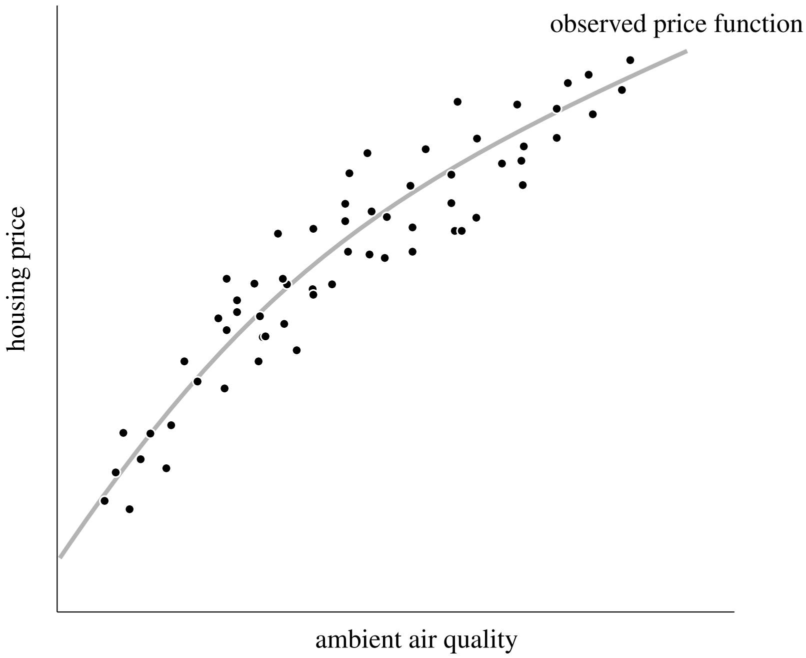

Figure 1: Housing price and ambient air quality

so that $p(q)$ is continuous and differentiable in q. Let us divide the entire goods into two broad categories: the differentiated good and everything else, the latter of which is assumed to be homogeneous (i.e., not differentiated). We consider an economy consisting of $N \in \mathbb{N}$ consumers. The problem of consumer i is formulated as

$$\max_{q_i, y_i} U^i(q_i, y_i) \quad \text{s.t.} \quad p(q_i) + y_i = m_i, \tag{1}$$

where y_i is the amount of numéraire consumed by this consumer. Denote by $q_i^d(p, m_i), y_i^d(p, m_i)$ the demand functions. It follows from the first-order condition that

$$p'(q_i^d(p, m_i)) = \frac{U_q^i(q_i^d(p, m_i), y_i^d(p, m_i))}{U_y^i(q_i^d(p, m_i), y_i^d(p, m_i))}. \tag{2}$$

What is meant by this equation can be seen in the top panel of Figure 2. The figure must look familiar, except that the budget line here is not represented by a straight line. This is because the price function is not linear in q. For each consumer, the marginal rate of substitution, U_q^i / U_y^i, is equalized with the slope p' of the price function. Hence, as in the standard utility maximization problem, consumer's indifference curve must be tangential to the budget line at the optimal consumption bundle. Just for illustrative purposes, Figure 2 depicts three utility maximization problems associated with three different consumers. Different consumers may have distinct preferences and therefore have different shapes of indifference curve. The shapes of the three budget lines are identical. The location of each line, however, shifts up and down, depending on the income level of each consumer.

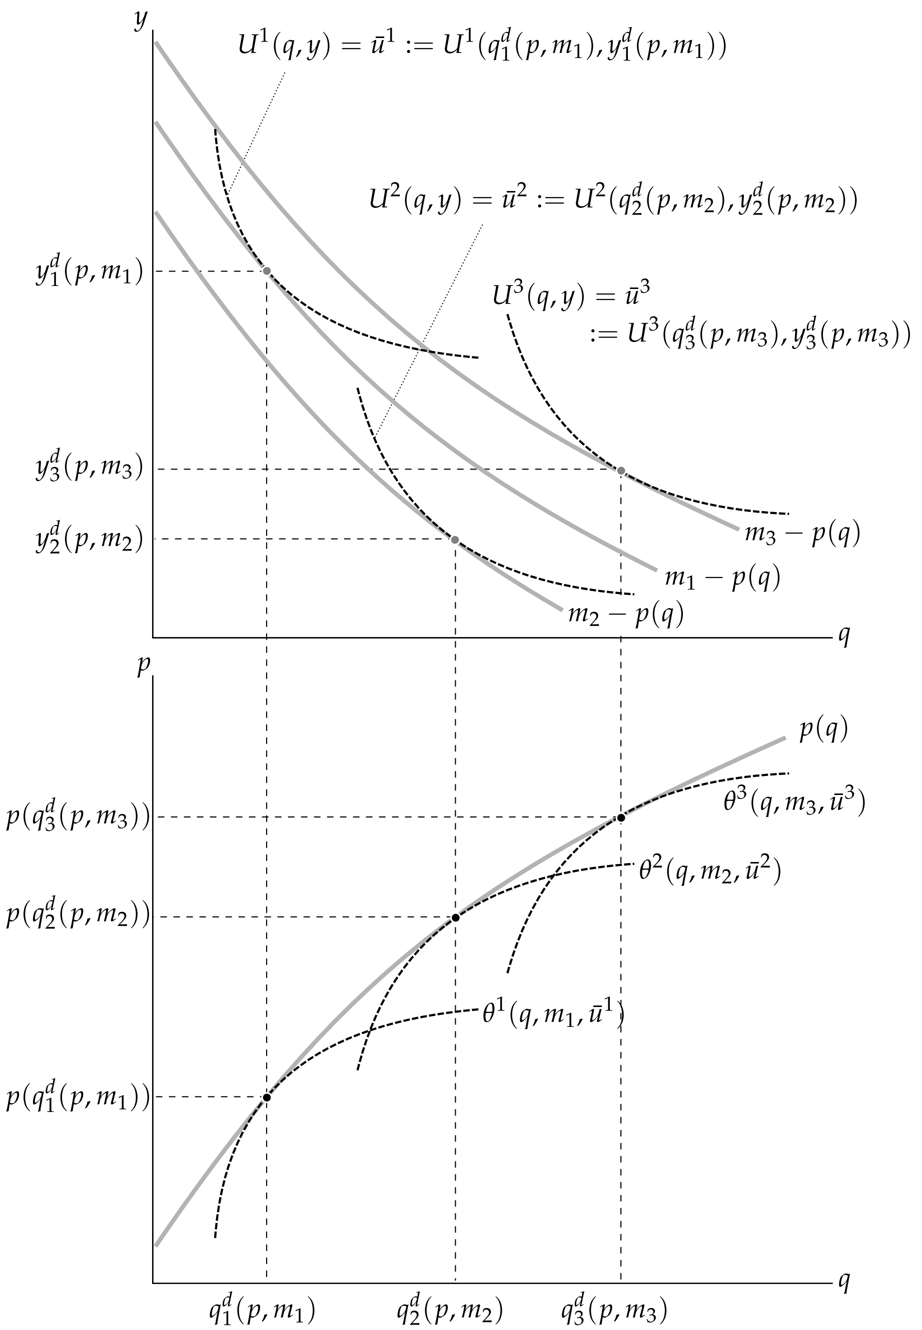

Figure 2: Consumer's choice, price function, and bid function

To better connect each consumer's preference with the price function, let us introduce the concept of *bid function*. A bid function represents the amount of money a consumer is willing to pay for alternative values of q at a given level of income and utility. More formally, a bid function of consumer i is implicitly defined by $\theta^i(q, m, u)$ such that

$$U^i(q, m - \theta^i(q, m, u)) = u \quad \forall (q, m, u). \tag{3}$$

Let $\bar{u}_i := U^i(q_i^d(p, m_i), y_i^d(p, m_i))$ be the maximized level of utility of consumer i. Then obviously,

$$U^i(q_i^d(p, m_i), m_i - p(q_i^d(p, m_i))) = \bar{u}_i \tag{4}$$

because $m_i - p(q_i^d(p, m_i)) = y_i^d(p, m_i)$. Comparing (3) and (4), we can see that the consumer i's bid function evaluated at $(q, m, u) = (q_i^d(p, m_i), m_i, \bar{u}_i)$ coincides with $p(q_i^d(p, m_i))$, namely,

$$\theta^i(q_i^d(p, m_i), m_i, \bar{u}_i) = p(q_i^d(p, m_i)). \tag{5}$$

On the other hand, by partially differentiating (3) with respect to q, we obtain

$$U_q^i(q, m - \theta^i(q, m, u)) - U_y^i(q, m - \theta^i(q, m, u))\theta_q^i(q, m, u) = 0, \tag{6}$$

or equivalently,

$$\theta_q^i(q, m, u) = \frac{U_q^i(q, m - \theta^i(q, m, u))}{U_y^i(q, m - \theta^i(q, m, u))}. \tag{7}$$

Combining (2), (5), and (7) yields

$$\theta_q^i(q_i^d(p, m_i), m_i, \bar{u}_i) = \frac{U_q^i(q_i^d(p, m_i), y_i^d(p, m_i))}{U_y^i(q_i^d(p, m_i), y_i^d(p, m_i))} = p'(q_i^d(p, m_i)), \tag{8}$$

meaning that the bid function $\theta^i(q, m_i, \bar{u}_i)$ is tangential to the price function $p(q)$ at $q = q_i^d(p, m_i)$. The bottom panel of Figure 2 illustrates this fact.

As you might have already noticed, the bid curves and the indifference curves are closely connected. In fact, the bid curves in the bottom panel of Figure 2 can be obtained simply by flipping the indifference curves in the top panel upside down. Therefore, these two curves contain exactly the same information about consumer's preference.

1.2 Producer's decision

So far we have taken the price function $p(q)$ given, putting our focus on the demand side of the market. One might wonder where $p(q)$ came from in the first place. As is the case in the standard model of competitive market, the price is determined so that the demand and the supply are balanced. This should become clear if we explicitly introduce into the model the supply side of the market.[1]

[1] The theoretical basis for the hedonic price function we discuss here was first developed by Rosen (1974).

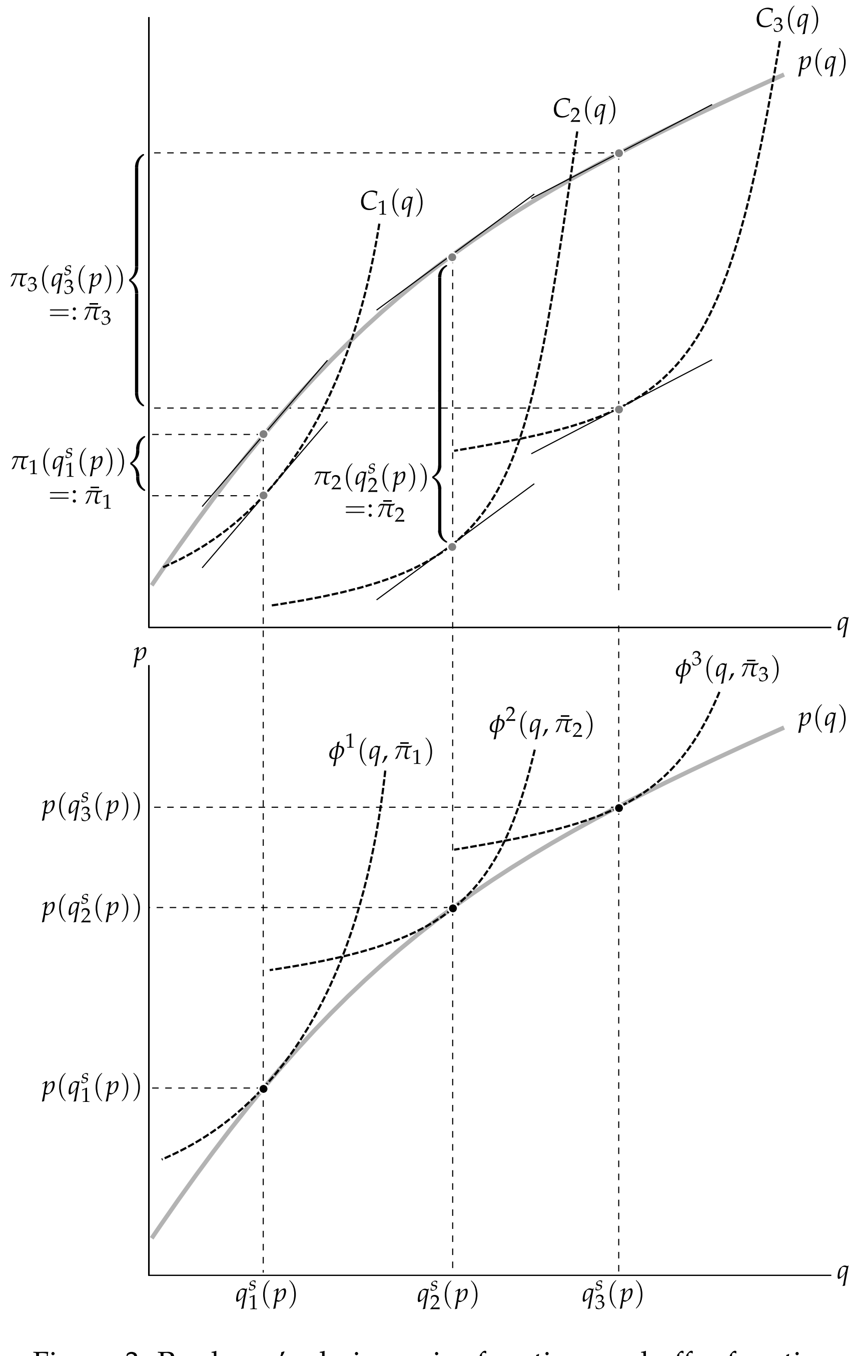

Figure 3: Producer's choice, price function, and offer function

For simplicity, assume that there are exactly N firms in the market and each of them produces only one differentiated good of a particular quality. Firm j's technology is represented by a cost function $C_j(q)$. The problem of firm j is then formulated as

$$\max_{q_j} \pi_j(q_j) := p(q_j) - C_j(q_j). \tag{9}$$

Let $q_j^s(p)$ be the supply function, which obviously depends on the price function p. Then the first-order condition implies that

$$\pi'(q_j^s(p)) = 0, \tag{10}$$

or equivalently,

$$p'(q_j^s(p)) = C_j'(q_j^s(p)). \tag{11}$$

This condition simply states that firms choose the quality of their product in such a way that the marginal revenue and the marginal cost are equalized. Unlike the standard profit maximization in the competitive market, the marginal revenue is not constant, reflecting the fact that the price function is not linear in q. See the top panel of Figure 3 for a diagramatic illustration.

To establish the link between firm's behavior and the price function, let us introduce the concept of *offer function*. An offer function represents the amount of money a firm is willing to accept for producing alternative values of q at a given level of profit. To be more formal, we define the offer function of firm j by $\phi^j(q, \pi)$ such that

$$\phi^j(q, \pi) - C_j(q) = \pi \quad \forall (q, \pi). \tag{12}$$

Let $\bar{\pi}_j := \pi_j(q_j^s(p))$ be the maximized profit of firm j. Then it should be easy to see

$$p(q_j^s(p)) - C_j(q_j^s) = \bar{\pi}_j. \tag{13}$$

Combining (12) with (13) yields

$$\phi^j(q_j^s(p), \bar{\pi}_j) = p(q_j^s(p)) \quad \text{and} \quad \phi_q^j(q_j^s(p), \bar{\pi}_j) = p'(q_j^s(p)), \tag{14}$$

which means that the offer function $\phi^j(q, \bar{\pi}_j)$ is tangential to the price function $p(q)$ at $q = q_j^s(p)$. This fact is depicted in the bottom panel of Figure 3.

In the supply side of the model, similar to the demand side, there exists a close connection between the offer curve and the cost function. To see this, notice that (12) implies

$$\phi^j(q, \pi) = C_j(q) + \pi. \tag{15}$$

This means that the offer curve is nothing but the cost function, except for its location along the vertical axis. One can obtain the offer function for each firm by just shifting their cost function upwards by the length of π. This should be clear if you compare the top and bottom panels of Figure 3. Hence, the offer function contains the information about firm's technology embodied in the cost function.

1.3 Equilibrium

Combining the demand side and the supply side of the model, we can now provide a theoretical underpinning of the hedonic price function observed in the market. Let us begin with the formal definition of equilibrium.

> Definition 1.1: **Equilibrium**
>
> Consider a market where differentiated goods with varying qualities are traded among $N \in \mathbb{N}$ consumers and N firms. We say that the market is in equilibrium if there exists a price function $p_* : \mathbb{R}_+ \to \mathbb{R}_{++}$ such that
>
> $$q_i^d(p_*, m_i) = q_{j(i)}^s(p_*) \quad \forall i \in \{1, 2, \ldots, N\} \tag{16}$$
>
> for some bijection $j : \{1, 2, \ldots, N\} \to \{1, 2, \ldots, N\}$, where $q_i^d(\cdot)$ is the demand function of consumer i and $q_{j(i)}^s(\cdot)$ is the supply function of firm $j(i)$.

The price function is eventually adjusted so that each consumer's demand can be matched up with each firm's supply. At equilibrium, there must exist N pairs of consumer and firm who choose the product of exactly the same quality. Geometrically, this means that for each pair of consumer and firm, the bid and offer curves touch each other at a particular location in the quality spectrum. Since bid and offer curves are both tangential to the price function, the equilibrium price function therefore lies in between these two curves. In other words, as shown in Figure 4, the price function is a bouble envelope of families of bid and supply curves.

Being a double envelope, the equilibrium price function contains two types of information: consumer's preference and firm's technology. At each point q where the differentiated good is actually traded between consumers and firms, we have

$$\theta_q^i(q, m_i, \bar{u}_i) = p_*'(q) = \phi_q^{j(i)}(q, \bar{\pi}_{j(i)}), \tag{17}$$

or equivalently,

$$\frac{U_q^i(q, m_i - p_*(q))}{U_y^i(q_i, m_i - p_*(q))} = p_*'(q) = C_{j(i)}'(q). \tag{18}$$

Hence, by examining the slope of the price function, we can infer the marginal rate of substitution of consumers and the marginal cost of firms. We should notice here that unlike the market of homogeneous goods, the marginal rates of substitution are not equalized among consumers. At each point q, the slope $p_*'(q)$ only represents the marginal rate of substitution of a particular consumer. The same thing can be said about firms, whose marginal costs are not equalized in the market of differentiated good.

Except that the hedonic price function should be increasing in q, little information about its functional form can be deduced from theory. The shape is jointly determined by the distribution of income and preferences among consumers and the distribution of technology among firms.

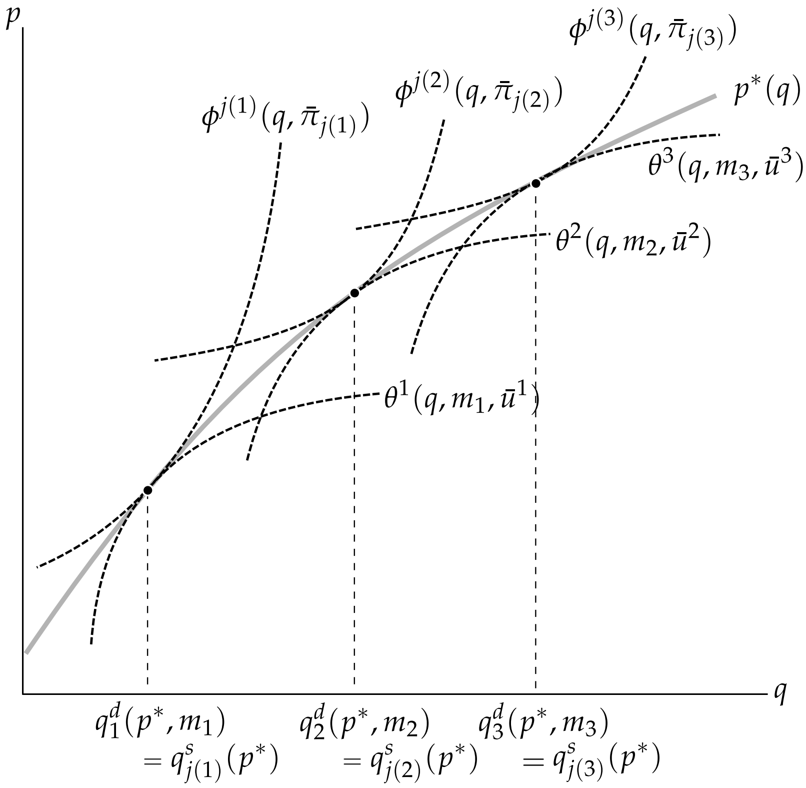

Figure 4: Equilibrium

2 Welfare measures

Consider a pubic policy which improves the air quality at a particular location in an otherwise heavily polluted region. If this policy is successfully implemented, those households living in the targeted location will experience an improvement in the ambient air quality. To make the exposition easier, let us assume that there is only one household living in the location. Then, the associated welfare change is measured in monetary term by how much this household is willing to pay for the expected improvement of ambient air quality. This in general requires information about the household's preference.

The discussion in the preceding section suggests that the hedonic price function contains information about consumer's preference and therefore may be used for valuing the welfare change due to provision of public goods. But how exactly can we compute the policy-induced welfare effect from the hedonic price function? Is it possible to construct an exact welfare measure solely based on observable data? In what follows, we will try to answer these questions based on the model of quality choice we discussed in the preceding section.

2.1 Case with no transaction cost

Consider an individual whose utility maximization problem is given by

$$\max_{q,y} U(q,y) \quad \text{s.t.} \quad p(q) + y = \bar{m}. \tag{19}$$

Let $(q^d(p,\bar{m}), y^d(p,\bar{m}))$ be the solution of this problem, which satisfies

$$p'(q^d(p,\bar{m})) = \frac{U_q(q^d(p,\bar{m}), y^d(p,\bar{m}))}{U_y(q^d(p,\bar{m}), y^d(p,\bar{m}))} \tag{20}$$

and

$$p(q^d(p,\bar{m})) + y^d(p,\bar{m}) = \bar{m}. \tag{21}$$

In other words, this individual buys a house with quality $q^d(p,\bar{m})$ and owns $y^d(p,\bar{m})$ units of numéraire. The associated level of utility is given by $V(p,\bar{m}) := U(q^d(p,\bar{m}), y^d(p,\bar{m}))$. Now suppose that the ambient air quality of the house improves by $\Delta q > 0$ and nothing else is affected. We want to measure the welfare impact of this improvement in monetary term.

Assume, for the moment, that there are a sufficiently large number of buyers and sellers in the market and the transaction cost is negligible. In this case, the owner of the house can easily sell their possessions (the house and the numéraire) in the market and then repurchase a more preferable pair of q and y under the renewed budget. Hence, once the ambient air quality improves, the owner solves the utility maximization problem

$$\max_{q,y} U(q,y) \quad \text{s.t.} \quad p(q) + y = \bar{m}', \tag{22}$$

where $\bar{m}'$ is the renewed budget, which is defined by

$$\bar{m}' := p(q^d(p, \bar{m}) + \Delta q) + y^d(p, \bar{m}). \tag{23}$$

Notice that $\bar{m}'$ is strictly larger than $\bar{m}$ because

$$\begin{aligned} \bar{m}' - \bar{m} &= \Big(p(q^d(p, \bar{m}) + \Delta q) + y^d(p, \bar{m})\Big) - \Big(p(q^d(p, \bar{m})) + y^d(p, \bar{m})\Big) \\ &= p(q^d(p, \bar{m}) + \Delta q) - p(q^d(p, \bar{m})) > 0, \end{aligned} \tag{24}$$

where the inequality follows from the fact that the hedonic price function p is increasing in q. Let $(q^d(p, \bar{m}'), y^d(p, \bar{m}'))$ be the solution of (22), which satisfies

$$p'(q^d(p, \bar{m}')) = \frac{U_q(q^d(p, \bar{m}'), y^d(p, \bar{m}'))}{U_y(q^d(p, \bar{m}'), y^d(p, \bar{m}'))} \tag{25}$$

and

$$p(q^d(p, \bar{m}')) + y^d(p, \bar{m}') = \bar{m}'. \tag{26}$$

The associated level of utility is given by $V(p, \bar{m}') := U(q^d(p, \bar{m}'), y^d(p, \bar{m}'))$.

We employ two concepts of welfare measure: *compensating variation* and *equivalent variation*. The compensating variation is the maximum amount of money that the owner of the house is willing to pay for improving q by the amount of Δq. In the present context, it is defined by $\overline{\mathrm{CV}}$ such that

$$V(p, \bar{m}' - \overline{\mathrm{CV}}) = V(p, \bar{m}), \tag{27}$$

which is equivalent to

$$\bar{m}' - \overline{\mathrm{CV}} = \bar{m}. \tag{28}$$

The equivalent variation, on the other hand, is the minimum amount of money that the same owner is willing to accept for giving up the opportunity to improve the level of q. In the present context, it is defined by $\overline{\mathrm{EV}}$ such that

$$V(p, \bar{m}') = V(p, \bar{m} + \overline{\mathrm{EV}}), \tag{29}$$

which is equivalent to

$$\bar{m}' = \bar{m} + \overline{\mathrm{EV}}. \tag{30}$$

It should be easy to see from (28) and (30) that

$$\overline{\mathrm{CV}} = \overline{\mathrm{EV}} = \bar{m}' - \bar{m}. \tag{31}$$

Combining (31) and (24), we obtain

$$\overline{\mathrm{CV}} = \overline{\mathrm{EV}} = p(q^d(p, \bar{m}) + \Delta q) - p(q^d(p, \bar{m})). \tag{32}$$

Therefore, in this case, the compensating variation and equivalent variation are identical and their value coincides with the price change of the house affected by the air-quality improvement.

What is particularly worth mentioning here is the fact that the right-hand side of (32) is entirely observable. No information about individual's preference is needed. All we need to know is the shape of the hedonic price function p and the degree Δq of quality improvement caused by a public policy. This is because in the absence of transaction cost, the entire welfare effect of quality improvement is capitalized in the asset value of the house owner.

2.2 Case with positive transaction cost

In reality, however, the cost of transaction and moving can be substantially large. Negotiating with potential buyers of your house could require a lot of work while finding a new location of your house could take a long time. Physically moving your furniture from one place to another is far from costless. If this cost is substantially high, you might find it reasonable to stay in the current location even if the ambient air quality of your house changes.

To illustrate how the welfare measures need to be modified in the presence of transaction cost, let us consider an extreme case where the transaction cost is prohibitively high. In such a case, individuals cannot sell their possessions in the market and hence are not allowed to adjust their behavior after the ambient air quality changes. The utility maximization problem is given by

$$\max_{q,y} U(q,y) \quad \text{s.t.} \quad p(q) + y = \bar{m}. \tag{33}$$

We denote by $(q^d(p,\bar{m}), y^d(p,\bar{m}))$ the solution of this problem. The utility of this individual is then characterized by $U(q^d(p,\bar{m}), y^d(p,\bar{m}))$. Suppose that the ambient air quality of the house improves by $\Delta q > 0$ and nothing else is affected. Then, since the transaction cost is sufficiently high, this individual necessarily chooses $(q^d(p,\bar{m}) + \Delta q, y^d(p,\bar{m}))$. The associated level of utility is given by $U(q^d(p,\bar{m}) + \Delta q, y^d(p,\bar{m}))$.

Recall that y is measured in units of money (i.e., numéraire). Thus, in the present context, it will be appropriate to define the compensating variation by the maximum amount of numéraire that the house owner is willing to pay for improving the ambient air quality. To be more formal, the compensating variation is defined by $\underline{\text{CV}}$ such that

$$U(q^d(p,\bar{m}) + \Delta q, y^d(p,\bar{m}) - \underline{\text{CV}}) = U(q^d(p,\bar{m}), y^d(p,\bar{m})). \tag{34}$$

Since $p(q^d(p,\bar{m})) + y^d(p,\bar{m}) = \bar{m}$, this may be rewritten as

$$U(q^d(p,\bar{m}) + \Delta q, \bar{m} - (p(q^d(p,\bar{m})) + \underline{\text{CV}})) = \bar{u}, \tag{35}$$

where

$$\bar{u} := U(q^d(p,\bar{m}), y^d(p,\bar{m})). \tag{36}$$

By the definition of bid function, (3), we then have

$$\theta(q^d(p,\bar{m}) + \Delta q, \bar{m}, \bar{u}) = p(q^d(p,\bar{m})) + \underline{\text{CV}} \tag{37}$$

or equivalently,

$$\begin{aligned} \underline{\text{CV}} &= \theta(q^d(p,\bar{m}) + \Delta q, \bar{m}, \bar{u}) - p(q^d(p,\bar{m})) \\ &= \theta(q^d(p,\bar{m}) + \Delta q, \bar{m}, \bar{u}) - \theta(q^d(p,\bar{m}), \bar{m}, \bar{u}), \end{aligned} \tag{38}$$

where the second line uses (5). Therefore, when the transaction cost is sufficiently high, the compensating variation can be simply expressed as how much the bid function changes in response to the air-quality improvement.

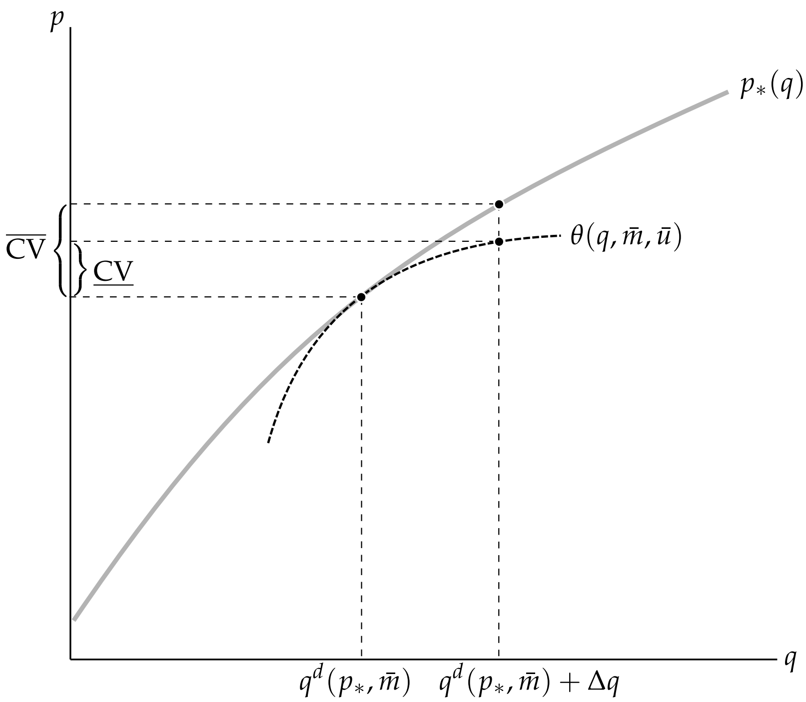

Figure 5: Upper and lower bounds of CV

Observe

$$\begin{aligned}\overline{\mathrm{CV}} - \underline{\mathrm{CV}} &= \Big(p(q^d(p,\bar{m}) + \Delta q) - p(q^d(p,\bar{m}))\Big) \\ &\qquad - \Big(\theta(q^d(p,\bar{m}) + \Delta q, \bar{m}, \bar{u}) - \theta(q^d(p,\bar{m}), \bar{m}, \bar{u})\Big) \\ &= p(q^d(p,\bar{m}) + \Delta q) - \theta(q^d(p,\bar{m}) + \Delta q, \bar{m}, \bar{u}), \end{aligned} \tag{39}$$

where the second equality follows from (5). Since the hedonic price curve is an envelope of the bid curves, the former curve always lies above the latter. This implies that

$$p(q^d(p,\bar{m}) + \Delta q) > \theta(q^d(p,\bar{m}) + \Delta q, \bar{m}, \bar{u}) \tag{40}$$

for any $\Delta q \neq 0$, as depicted in Figure 5. Combining (39) with (40) hence yields

$$\overline{\mathrm{CV}} > \underline{\mathrm{CV}}. \tag{41}$$

This should come as no surprise given the fact $\overline{\mathrm{CV}}$ completely ignores the transaction cost while $\underline{\mathrm{CV}}$ assumes the infinite amount of transaction cost. The value of air-quality improvement should be lower if costly adjustments are required to fully enjoy the benefit of the improvement.

In a more realistic situation, the transaction cost is likely to be neither negligible nor prohibitive. Then the compensating variation should lie somewhere in between $\underline{\mathrm{CV}}$ and $\overline{\mathrm{CV}}$. Put differently, $\underline{\mathrm{CV}}$ provides a lower bound of CV while $\overline{\mathrm{CV}}$ provides an upper bound.

2.3 Empirical application

A recent study by Chay and Greenstone (2005) nicely illustrates the usefulness of hedonic price model. In 1971, the US congress passed the Clean Air Act Amendments, which established a federal standard for local air pollution concentrations. The air pollutant in question then was tiny airborne particles known as total suspended particulate matters or TSPs. When the federal TSPs standard was violated in a county, every major source of pollution locating in the county was required to make a substantial investment in pollution reduction activities. Largely due to this regulation, the TSPs levels in the regulated counties dropped on average by 10 μg/m^3 in the period of 1970 to 1980. During the same period, the housing price surged in those countries. Then the question is how much of the observed increase in property value is attributed to the improvement of local air quality.

Based on detailed data on county-level air pollution, county characteristics, and property values, they estimated the hedonic price function as

$$p = e^{\beta_0 + \beta_x x + \beta_q q + \varepsilon} \tag{42}$$

or equivalently,

$$\ln(p) = \beta_0 + \beta_x x + \beta_q q + \varepsilon, \tag{43}$$

where p is the housing price, q is the level of TSPs (μg/m^3), and x is a vector of other relevant covariates. The estimated value of β_q is around -0.28, which means that 1 μg/m^3 decline in TSPs concentration will raise the average property value by 0.28 percent. Hence, during 1970 to 1980, it is estimated that the average property value increased by $0.28 \times 10 = 2.8$ percent due to the improvement of local air quality in the regulated counties. Since the average value of a house in those counties was 86,900 dollars in 1970, and since 19 million houses existed there, a rough calculation leads us to conclude that the aggregate benefit associated with the air-quality improvement is $86,900 \times 0.028 \times 19,000,000 = 46$ billion dollars in total.

3 Value of a statistical life

The hedonic price model we learned in the preceding section plays an important role in evaluating a public policy, especially when the primary benefit of the policy is a reduction in mortality risks. Measuring the benefit associated with mortality-risk reduction can be critical in cost-benefit analysis. The quantified benefits of clean air policies or drinking water regulations, for instance, are typically dominated by the welfare gain from saving those lives that would otherwise be lost due to air or water pollution. In fact, it is estimated that more than 80% of the benefits of the Clean Air Act in the United States is accounted for by the reductions in mortality risks.[2]

[2] According to USEPA (2011), the economic benefits of the Clean Air Act Amendments will be 2 trillion US dollars in 2020 while the corresponding cost is only 65 billion dollars. About 85% of the benefits are attributed to reductions in premature mortality associated with ambient particulate matter.

The benefit of mortality-risk reduction is often expressed in terms of *value of a statistical life* or VSL for short. VSL is best interpreted as the aggregate willingness to pay for marginally reducing the risk of death. In order to get a sense of the concept, let us consider a simple hypothetical example. Suppose that in a country populated by 1 million people, the government is introducing a new food safety regulation. This regulation is expected to reduce the annual risk of death caused by a food-related illness by 0.000001. This means that the food safety regulation, once effectively implemented, will save 1 person annually in the country as a whole. In other words, one *statistical* life is saved by this regulation. The life is saved only in a statistical sense because the identity of the saved person cannot be specified ex-ante. Suppose that each person in this country is willing to pay 4 dollars a year for reducing mortality risk by 0.000001. Then, the aggregate willingness to pay for saving one statistical life — the value of a statistical life — is $4 \times 1000000 = 4$ million dollars.[3]

Theoretically speaking, VSL is defined as the average marginal rate of substitution between money and mortality risk in a predetermined time period. Consider an economy consisting of $N \in \mathbb{N}$ consumers. Suppose that consumer i's preference is represented by a utility function $U^i(\pi_i, y_i)$, where $\pi_i \in [0, 1]$ is the probability of death during a fixed period of time (say, a year) and y_i is the amount of numéraire consumed by this individual. It is reasonable to assume that U^i is decreasing in π_i and is increasing in y_i. Then, for a given $\Delta\pi \in \mathbb{R}_{++}$, there must exist $\mathrm{CV}_i \in \mathbb{R}_{++}$ such that

$$U^i(\pi_i - \Delta\pi, y_i - \mathrm{CV}_i) = U^i(\pi_i, y_i). \tag{44}$$

This CV_i is consumer i's willingness to pay for reducing mortality risk by $\Delta\pi$. Notice that if $\Delta\pi$ (and hence CV_i) is sufficiently small, the left-hand side of this equation may be approximated as

$$U^i(\pi_i - \Delta\pi, y_i - \mathrm{CV}_i) \approx U^i(\pi_i, y_i) - U^i_\pi(\pi_i, y_i)\Delta\pi - U^i_y(\pi_i, y_i)\mathrm{CV}_i. \tag{45}$$

Combining (44) with (45) yields

$$\mathrm{CV}_i \approx -\frac{U^i_\pi(\pi_i, y_i)}{U^i_y(\pi_i, y_i)}\Delta\pi \tag{46}$$

for sufficiently small $\Delta\pi$. Consider a public policy which can save one statistical life in this economy. This means that the mortality risk can be reduced by $\Delta\pi := 1/N$ for every individual in the economy. Then, the aggregate willingness to pay for such a policy (i.e., VSL) is

$$\mathrm{VSL} = \sum_{i=1}^{N} \mathrm{CV}_i \approx -\frac{1}{N} \sum_{i=1}^{N} \frac{U^i_\pi(\pi_i, y_i)}{U^i_y(\pi_i, y_i)}, \tag{47}$$

the right-hand side of which is the average marginal rate of substitution between risk and money (numéraire good). The approximation is fairly good as long as N is sufficiently large.

[3]More generally, one can imagine a public policy which could save $n \in \mathbb{N}$ statistical lives. Then the associated benefit is computed as $n \times$ VSL.

3.1 Hedonic wage model

To construct a reliable estimate of VSL, one needs to find a market where people's preference for risk reduction is revealed. To be more precise, we need to know the marginal rate of substitution between risk and money. Since there is no market where the probability of death is directly traded, we usually turn to the existing markets — labor markets, typically — which implicitly involve tradeoffs between risk and money. The idea is simple. If we can observe how much additional money is required to compensate someone for having a higher mortality risk on the job, that would reveal how individuals trade off safety for money.

Consider, for instance, a labor market where two mining companies are seeking workers. The job opportunities offered by these companies are virtually identical, except for on-the-job risk of fatal accidents. One of the companies is operating in a dangerous region with a higher frequency of life-threatening accidents while the other in a relatively safer region. If the wages are the same between the two jobs, most people would prefer working in the safer region. Getting people to take the job in the dangerous region hence requires some incentive, such as higher pay. The difference in wages for the dangerous and safe jobs, which is sometimes called the *compensating wage differential*, tells us how much people need to be paid to accept the risk. This in turn reveals how much people are willing to pay for reducing the risk.

The idea can be formalized by a straightforward extension of the hedonic price model. Consider a labor market consisting of $N \in \mathbb{N}$ individuals. We assume that a large number of job opportunities are offered in the market with varying degrees of on-the-job mortality risk and wage. For simplicity, we assume that other job-related characteristics are fully controlled for. In other words, all the job opportunities in the labor market are identical except for the on-the-job mortality rate and wage rate. The problem of individual i is given by

$$\max_{\pi_i, y_i} U^i(\pi_i, y_i) \quad \text{s.t.} \quad y_i = m_i + w(\pi_i), \tag{48}$$

where $w(\pi)$ is the *hedonic wage function*, which maps a wage rate w to a particular level of on-the-job mortality risk π. If the market is sufficiently competitive, a job with a higher mortality risk needs to be accompanied by a higher wage rate. Hence, it is reasonable to assume that $w(\pi)$ is increasing in π.

Denote by $\pi_i^d(w, m), y_i^d(w, m)$ the demand functions. It follows from the first-order condition that

$$w'(\pi_i^d(w, m_i)) = -\frac{U_\pi^i(\pi_i^d(w, m_i), y_i^d(w, m_i))}{U_y^i(\pi_i^d(w, m_i), y_i^d(w, m_i))}. \tag{49}$$

This expression already reveals that the slope of the hedonic wage function contains the information required for computing VSL. In fact, using (47) and (49), we obtain

$$\text{VSL} \approx \frac{1}{N} \sum_{i=1}^{N} w'(\pi_i^d(w, m_i)) \tag{50}$$

for sufficiently large $N \in \mathbb{N}$.

3.2 Empirical application

The value of a statistical life has been estimated by a large number of studies. Many of the recent studies are based on a data set called the Census of Fatal Occupational Injuries (CFOI), which is prepared by the US Bureau of Labor Statistics. Table 1 presents a list of on-the-job incidences of death by occupation and industry, which is computed based on CFOI. The table has 10 occupational groups and divides the industries into 9 categories. The overall probability of death implied by the table is 0.00004, but the numbers significantly vary across occupation groups and industries. In Finance industry, for instance, only 1.36 fatal incidence occurs per 100,000 workers while the rate sharply increases up to 25.99 in Mining industry.

Using this data, Viscusi (2004) estimated the hedonic wage function

$$\ln(w) = \beta_0 + \beta_x x + \beta_\pi \pi + \varepsilon, \tag{51}$$

where w is worker's hourly wage rate, x is a vector of personal and job characteristics, and π is the occupation- and industry-specific annual probability of death. Notice that we may write

$$\beta_\pi = \frac{\partial \ln(w)}{\partial \pi} = \frac{1}{w}\frac{\partial w}{\partial \pi} \tag{52}$$

or

$$\frac{\partial w}{\partial \pi} = w\beta_\pi. \tag{53}$$

The average wage rate is $w = 13.92$ and the estimated value of β_π is 170 for $\Delta\pi = 0.00001$. Hence,

$$\frac{\partial w}{\partial \pi}\Delta\pi = w\beta_\pi\Delta\pi = 13.19 \times 170 \times 0.00001 = 0.02366, \tag{54}$$

which means that the hourly wage rate increases by 0.02366 dollars if the annual on-the-job probability of death increases by 0.00001. In other words, people are willing to give up 0.02366 dollars per working hour if the risk of fatal accident can be reduced by 0.00001. Assuming that an average employee works for 2000 hours per year, we may compute the annualized value of the willingness to pay (or CV_i) as

$$\mathrm{CV}_i = 2000 \times \frac{\partial w}{\partial \pi}\Delta\pi = 2000 \times 0.02366 = 47.328. \tag{55}$$

To convert this value into the value of a statistical life, we need to aggregate over population of size $N = 1/\Delta\pi = 100000$. Namely,

$$\mathrm{VSL} = \sum_{i=1}^{N} \mathrm{CV}_i = \sum_{i=1}^{N} 47.328 = 4732800, \tag{56}$$

implying that the estimated value of a statistical life is about 4.7 million dollars.

Table 1: Annual incidence of fatality by occupation and industry (1992–97 average, per 100,000 employees)

	Mining	Constr.	Manuf.	Transp.	Wholesale	Retail	Finance	Services	Public	Total
Executive	5.80	4.89	1.84	2.22	3.34	4.77	1.67	1.56	2.60	2.38
Professional	6.86	2.64	1.20	2.72	2.88	1.80	0.71	1.13	2.72	1.30
Technicians	10.67	15.19	2.49	21.03	3.82	0.87	0.77	1.74	7.88	3.92
Sales	5.00	4.86	3.54	2.23	3.26	3.87	1.45	2.11	2.60	3.30
Administrative support	0.51	0.98	0.56	1.41	0.63	0.59	0.44	0.47	0.97	0.66
Service	22.22	4.76	5.66	6.06	5.45	1.69	5.21	1.90	11.26	2.92
Precision production	38.54	11.38	3.63	7.48	7.82	3.11	3.13	5.43	11.58	7.59
Machine operators	24.31	30.41	2.15	6.64	9.90	1.43	4.17	2.33	14.67	2.81
Transportation	42.90	20.88	15.79	28.82	14.97	11.86	10.61	12.02	25.89	21.47
Handlers & cleaners	45.83	31.41	7.57	12.93	10.09	3.60	12.35	10.40	42.65	12.02
Total	25.99	12.62	3.02	10.75	5.19	3.29	1.36	1.76	5.72	4.02

Source: Taken from Table 1 of Viscusi (2004), which is based on the Census of Fatal Occupational Injuries prepared by the US Bureau of Labor Statistics.

References

Chay, K.Y. and M. Greenstone (2005), "Does air quality matter? Evidence from the housing market," *Journal of Political Economy*, 113(2):376–424.

Rosen, S. (1974), "Hedonic pries and implicit markets: product differentiation in pure competition," *Journal of Political Economy*, 82(1):34–55.

USEPA (2011), *Health and Welfare Benefits Analyses to Support the Second Section 812 Benefit-Cost Analysis of the Clean Air Act*, United States Environmental Protection Agency.

Viscusi, W.K. (2004), "The value of life: estimates with risks by occupation and industry," *Economic Inquiry*, 42(1):29–48.