Week 3

Market Failure Due to Information Asymmetry
Adverse Selection and Signalling

Information asymmetry refers to the fact that the buyer and the seller of a commodity may have different amounts of information about that commodity's attributes.

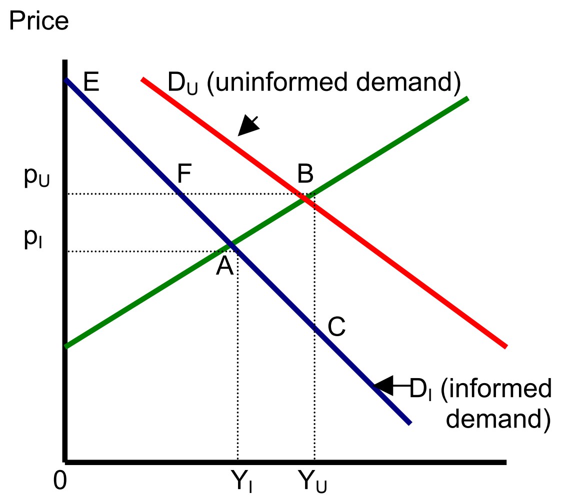

Figure 10:
Market Failure From Information Asymmetry

Source: Weimer and Vining (1999), p. 108

In Figure 10, D_U and D_I represent the consumer's demand schedule in, respectively, the absence and presence of perfect information about its quality: they are, respectively, the consumer's 'uninformed' and 'informed' demands[1]. The quantity *actually* purchased is Y_U and this is greater than Y_I, the quantity the consumer *would have* bought had he been fully informed.

The gain in producer's surplus from 'over-consumption' is p_uBAp_I. The loss in consumer's surplus from over-consumption is: Ap_IE – (Op_UBY_U-$OECY_U$) = Ap_IE – (p_UEF – FBC) = Ap_Ip_UF+FBC. So, the net loss to society from 'over-consumption' is loss in consumer's surplus less gain in producer's surplus = Ap_Ip_UF+(AFB+ABC) – (Ap_Ip_UF+AFB)=**ABC**.

[1] See Peltzman (1973) for the basic analysis and McGuire, Nelson and Spavins (1975) for a discussion of the empirical problems in using this approach.

When consumers overestimate quality, through a lack of information, producers lack incentives to provide information. Accurate information would lead to a lower surplus for the producer. An analysis, identical to that above, would apply if, due to lack of information, consumers underestimated quality so that there was 'under-consumption'. Now, however, producers would have an incentive to provide information since accurate information would now lead to a higher surplus for the producer.

Commodities for which information is required for satisfactory consumption may be divided into *search* goods and *experience* goods (Nelson, 1970). Information about the attributes of a search good can be determined *prior* to purchase (for example, how comfortable a sofa in a shop is likely to be) whereas information about an experience good can only be obtained *after* purchase (the quality of food in a new restaurant; the reliability of a second-hand car). The effectiveness of an information-gathering strategy depends upon:

(i) the variance in the quality of the good
(ii) the frequency of purchase
(iii) the full price of the good, including any harm from use
(iv) the cost of searching

Search Goods

Search goods may be thought of as a sampling process in which a consumer pays a cost of *$s* to inspect a particular price-quantity combination of a good. The good is rejected if price exceeds the consumer's marginal value for the good. Then the consumer either pays another *$s* to sample another price-quantity combination or stops searching. If the marginal valuation exceeds price the consumer either makes a purchase or continues to search in the hope of finding a more favourable surplus.

The greater the heterogeneity in quality and/or the higher the search costs, the greater the potential for inefficiency through information asymmetry. The point is that search goods rarely involve information asymmetry that lead to significant and persistent inefficiency calling for public policy intervention.

Experience Goods

With experience goods, consumers have bear the search cost and the *full price* (p^*) of a good in order to learn about its quality. The full price of a good (Oi, 1973) may be defined as follows. Suppose that a consumer buys Y units at a price of p per unit and that the probability of a defective item is $1-q$; then, on average, the consumer expects $Z=Yq$ units. If a bad unit inflicts a damage of W then the total cost of the purchase (C), and the full price (p^*), are defined as:

$$C = pY + W(Y - Z) \text{ and } p^* = \frac{C}{Z} = \frac{p}{q} + W\frac{1-q}{q} \qquad (1)$$

Since the consumer has to purchase the good prior to discovering its quality one would expect that:

(a) sampling would be less frequent, the more expensive the good

(b) sampling would be less frequent, the more durable the good

Furthermore, the consumer may discover, after purchase, that the marginal value is less than price and may regret the purchase.

The Market for 'Lemons': Diagrammatic Analysis

A particular example of 'experience goods' is the market for used cars (Akerlof, 1970). There are two kinds of used cars being sold on the market: 'low-quality' and 'high-quality'. If both sellers and buyers knew whether a given car was low or high quality, there would be a market for low quality cars and a separate market for high quality cars (Figure 11, below).

Source: Pindyck and Rubinfeld (2001), p. 597

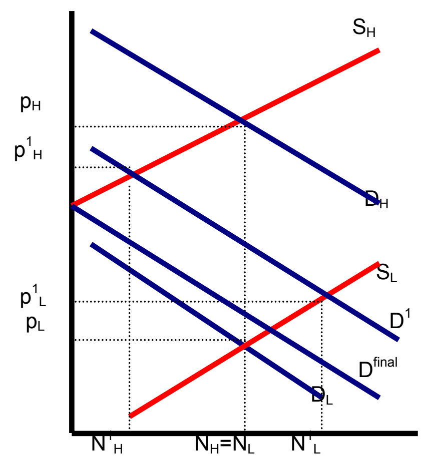

Figure 11
The Market for Lemons

In Figure 11, the price of high-quality cars is p_H and N_H of such cars are sold; the price of low-quality cars is p_L and N_L of such cars are sold. The demand and supply curves of high-quality cars (D_H and S_H) are above the demand and supply curves of low-quality cars (D_L and S_L). The number of high- and low-quality cars is the same (N_H=N_L), but p_H>p_L.

Now suppose that buyers cannot distinguish between high- and low-quality cars. So, if *N* cars were on the market, buyers would regard a given car to be as likely to be a low-quality as a high-quality car. So buyers would be prepared to pay: $p^1 = 0.5p^H + 0.5p^L$ for a car so the new demand curve D^1 lies half-way between D_H and D_L. The price of high-quality cars falls from p_H to p^1_H and the number of high-quality cars sold from N_H to i; and the price of low-quality cars rises from p_L to p^1_L and the number of low-quality cars rises from N_L to N^1_L. This causes buyers to revise downwards the chances of being offered a high-quality car – and to revise upwards the chances of being offered a low-quality car – causing a further leftward shift in the demand curve. The demand curve continues to shift until only low-quality cars are sold (D^{final} in Figure 11).

The Market for 'Lemons': Formal Analysis

Suppose, that the quality of a used car can be indexed by $q,\ q \in [0,1]$. If q is uniformly distributed over the closed interval *[0,1]*, then *E(q)=0.5*. Suppose that there are: a large number of buyers who are prepared to pay a price of αq ($\alpha \geq 1$), and a large number of sellers who willing to accept a price of q, for a car of quality q. If quality was observable, then a car of quality *q* would sell for some price: $p(q) \in (\alpha q, q)$.

But, if quality was not observable, then consumers would estimate the quality of a car by the *average* quality of cars offered on the market. This average quality, denoted $\overline{q}$, can be observed and the consumers' willingness to pay for a car is $\alpha\overline{q}$. Under this circumstance, suppose that the equilibrium price is $p>0$. Then, *only* sellers whose used car is of quality $q \leq p$ will offer their cars for sale, since for the other sellers *p* is less than their reservation price, *q*.

Since quality is uniformly distributed over the interval *[0,p]*, average quality will fall to $\overline{q} = p/2 < 0.5$. Consequently, buyers would only be prepared to pay $\alpha\overline{q} = \alpha(p/2) = (\alpha/2)p < p$ for a car. Hence, no cars would be sold at the price *p*. Since the price *p* was chosen arbitrarily, no used cars will be sold at any positive price *p>0.* Hence, the only equilibrium price is *p=0*, when the demand and supply of used cars is zero: *asymmetric information destroys the market for used cars!* [2]

Adverse Selection

Adverse selection arises when products of different qualities are sold at the same price because, prior to purchase, the buyer cannot distinguish between products of different qualities. Alternatively, adverse selection could arise because a product of uniform quality is sold at the same price to buyers of different qualities and, prior to sale, the seller cannot distinguish between consumers of different qualities[3]. Whatever the sources of adverse selection, the consequence is the same: low-quality products, or high-risk buyers, 'crowd out' high-quality products, or low-risk buyers, so that what is observed is an *adverse selection* of products (as sellers of high-quality products withhold their product) or an *adverse selection* of buyers (as low-risk customers withhold their custom).

Adverse selection represents market failure since 'good' products and 'good' customers are under-represented, and 'bad' products and 'bad' customers are over-represented, in the market. The source of the market failure is the *externality* between products and between customers: when a seller of a low-quality product increases sales, he lowers the average quality of the product on the market, reduces the price the consumer is willing to pay and, thereby, hurts sellers of high-quality products; when a high-risk person buys insurance, he raises the average risk of the contingency; this increases the average

---

[2] The analysis is from Varian (1992), p. 468.

[3] An example is the insurance industry where the same premium, for a policy against a particular contingency, is charged to different individuals embodying different levels of risk in respect of the insured contingency.

premium the insurance company charges and, thereby, hurts low-risk persons.

Under adverse selection, therefore, sellers of high-quality products will have an incentive to *signal* to the consumer the quality of their product. This may take the form of: *reputation*; *standardisation*; *informative advertising*; offering *warranties* in the event of defects. The signal may be offered through third parties: *recommendations* by friends or by consumer reports; *certification of quality* by a professional association. Educational qualifications, analysed below, are an important way that potential employees signal their worker-qualities to employers.

Education as a Market Signal

A model of the education market is due to Spence (1974). In this model, there are two types of workers: 'good' workers and 'bad' workers. Good workers have a marginal product of a_G and bad workers have a marginal product of a_B: $a_G > a_B$. A fraction θ of the workers are 'good', the remainder, $1-\theta$ are 'bad'. The production function is linear, so that if L_G good, and L_B bad, workers are employed, output is:

$$Y = a_G L_G + a_B L_B \tag{2}$$

If worker quality was easily observable, the wage paid to each group would equal its marginal product: $w_G = a_G$ and $w_B = a_B$. But if a firm cannot observe worker quality, it offers the average wage to each group:

$$w = \theta a_G + (1-\theta) a_B \tag{3}$$

Now suppose that workers can acquire education and that the cost of acquiring education is lower for good workers than for bad workers: Ω_G and Ω_B are the 'amounts' of education acquired by good and bad workers and π_G and π_B are the costs of one unit of education for good and bad workers, $\pi_G < \pi_B$. Then the total cost of education of good and bad workers is:

$$C_G = \pi_G \Omega_G \text{ and } C_B = \pi_B \Omega_B \tag{4}$$

There are now two decisions to be made:

(i) Workers have to decide how much education to acquire

(ii) Firms have to decide how much to pay workers with different levels of education

Assume that education does nothing to increase productivity; its only value is as a signal. Now the firm adopts the following decision rule: for an education level, Ω^*, pay a wage of a_G if $\Omega \geq \Omega^*$ and pay a wage of a_B if $\Omega < \Omega^*$. In other words, education is taken as an indicator of worker quality and Ω^* separates workers into good and bad workers.

If under this rule, good workers acquire a level of education Ω^* or more, and bad workers acquire a level of education less than Ω^*, then the education level of a worker will perfectly signal his quality. The question is: would it be worthwhile for a bad worker to acquire an education level Ω^*? The cost of doing so is $\pi_B\Omega^*$ and the benefit from doing so is the increase in wages: a_G-a_B. So a bad worker will not acquire Ω^* education if:

$$\pi_B\Omega^* > a_G - a_B \tag{5}$$

and a good worker will acquire Ω^* education if:

$$\pi_G\Omega^* < a_G - a_B \tag{6}$$

So provided Ω^* satisfied the condition:

$$\frac{a_G - a_B}{\pi_B} < \Omega^* < \frac{a_G - a_B}{\pi_G} \tag{7}$$

the education of a worker will perfectly signal his quality. This type of equilibrium is called a *separating equilibrium* since it allows each type of worker to make a choice which separates him from the other type.

If, however, $\pi_B\Omega^* < a_G - a_B$ bad workers will also acquire the education level Ω^* and if $\pi_G\Omega^* > a_G - a_B$ even good workers will not acquire any education. So $\Omega^* < \frac{a_G - a_B}{\pi_B}$ or $\Omega^* > \frac{a_G - a_B}{\pi_G}$ will lead to a *pooling equilibrium* in which both types of workers make the same choice and the firm has to pay the average wage *w* of equation (7).

The separating equilibrium is socially inefficient because each good worker pays to acquire the education level Ω^*, even though it does nothing to increase his productivity, simply to distinguish himself from a bad worker. Exactly the same output is produced with signalling as without signalling (equation (6)), it is just that the distribution of rewards is different. So, under the terms of the model, investment in education confers a private gain (to the good workers who can earn more than bad workers) but no social benefit.

Numerical Example

This example is from:
http://courses.temple.edu/economics/Econ_92/Game_Lectures/11th-Lemons/market_for_lemons.htm

Table 1

		Valuation	Repair Cost	Net Value
Bad Car	Buyer	**$3200**	**$1700**	**$1500**
	Seller	**$2700**	**$1700**	**$1000**
Good Car	Buyer	**$3200**	**$200**	**$3000**
	Seller	**$2700**	**$200**	**$2500**

If the seller can truthfully offer buyer a bad car they will strike a deal between $1500, $1000
If the seller can truthfully offer buyer a good car they will strike a deal between $3000, $2500

Seller can offer the buyer a warranty: under the terms of the warranty he offers to pay all repair costs associated with the car

Table 2

		Buyer	**Seller**
Bad Car	**Warranty**	**If p<$2700, G=0 Price is below seller's 'reservation' price**	**If p<$2700, G=$1000 Seller holds on to car***
		If p≥$2700, G=$3200-p	**If p≥$2700, G=p-$1700**
	No Warranty	**If p<$1000, G=0 Price is below seller's 'reservation' price**	**If p<$1000, G=$1000 Seller holds on to car**
		If p≥$1000, G=$3200- p - $1700	**If p≥$1000, G=p**
Good Car	**Warranty**	**If p<$2700, G=0 Price is below seller's 'reservation' price**	**If p<$2700, G=$2500 Seller holds on to car****
		If p≥$2700, G=$3200- p	**If p≥$2700, G=p - $200**
	No Warranty	**If p<$2500, G=0 Price is below seller's 'reservation' price**	**If p<$2500, G=$2500 Seller holds on to car**
		If p≥$2200, G=$3200- p - $200	**If p≥$2700, G=p**

* If p<$2700, he will, after paying the repair cost of $1700, be left with less than $1000 which is his reservation price for a bad car
** If p<$2700, he will, after paying the repair cost of $200, be left with less than $2500 which is his reservation price for a good car

Table 3

Price $	Bad Car				Good Car			
	Warranty		No Warranty		Warranty		No Warranty	
	Buyer	Seller	Buyer	Seller	Buyer	Seller	Buyer	Seller
1000	0	1000	500	1000	0	2500	0	2500
1200	0	1000	300	1200	0	2500	0	2500
1400	0	1000	100	1400	0	2500	0	2500
1500	0	1000	0	1500	0	2500	0	2500
1600	0	1000	-100	1600	0	2500	0	2500
1800	0	1000	-300	1800	0	2500	0	2500
2000	0	1000	-500	2000	0	2500	0	2500
2200	0	1000	-700	2200	0	2500	0	2500
2400	0	1000	-900	2400	0	2500	0	2500
2500	0	1000	-1000	2500	0	2500	500	2500
2600	0	1000	-1100	2600	0	2500	400	2600
2700	500	1000	-1200	2700	500	2500	300	2700
2800	400	1100	-1300	2800	400	2600	200	2800
2900	300	1200	-1400	2900	300	2700	100	2900
3000	200	1300	-1500	3000	200	2800	0	3000
3100	100	1400	-1600	3100	100	2900	-100	3100
3200	0	1500	-1700	3200	0	3000	-200	3200

Bad Car with Warranty: Net value of trade is positive for both parties for p≥2700
Bad Car without Warranty: Net value of trade is positive for both parties for p≤1500
Good Car with Warranty: Net value of trade is positive for both parties for p≥2700
Good Car without Warranty: Net value of trade is positive for both parties for 2500≤p≤3000

Now we make the assumption that buyer cannot distinguish between a good and a bad car: the maximum he is willing to pay for a bad car is $1500 and the maximum he is willing to pay for a good car is $3000. If he picks a car at random, there is an equal chance of getting a good and a bad car. So buyer will offer to pay p=0.5×1500+0.5×3000=$2250 for a randomly chosen car.

If the buyer has picked a bad car and offers $2250, seller will *accept*; but if the buyer has picked a good car and offers $2250, seller will *decline*.

So buyer knows that for $2250 he can never get a good car but only a bad car. So, with this knowledge, the maximum price he would offer is $1500 for a bad car. No good cars will be sold.

Suppose the seller offers a warranty with the good car, but not with the bad car. In other words, whether or not a warranty is being offered *signals* the quality of the car.

If no warranty is being offered, the buyer knows it is a bad car and he will offer $1000; If a warranty is being offered, the buyer knows it is a good car and he will offer $2700. In both cases he is offering the seller's reservation price.

Suppose the seller offers a warranty on both types of cars. Then he would receive $2700 for the car which would leave him $1000 after paying repair costs. So he has no incentive to offer warranty on a bad car. He will not remove the warranty from the good car since then he will receive an offer of $1000 for it from the buyer who cannot tell the difference between a good car and a bad car.

So the only rational course is for the seller to offer an warranty on the good cars but not on the bad cars. So before the warranty was offered, the buyer believed Pr(bad car)=0.5. This is his prior probability. But when a warranty is offered, this gives him further information: Pr(car is bad | warranty)=0; Pr(car is good | warranty)=1. These are his posterior probabilities and they establish a *separating equilibrium* by distinguishing between the two types of cars, depending upon whether or not they offer a warranty.

Bayes' Theorem
Reverend Thomas Bayes – an 18th century Presbyterian minister – proved what is, arguably, the most important theorem in statistics (see "In Praise of Bayes", The *Economist*, 28 September 2000).
Let T denote "theory" and D denote "data". Then the probability of the theorem being true, *given that the data has been observed*, is:

$$P(T \mid D) = \frac{P(T \cap D)}{P(D)} = \frac{P(D \mid T)P(T)}{P(D)} \tag{8}$$

where: $P(D) = P(D \mid T)P(T) + P(D \mid \bar{T})P(\bar{T})$, $\bar{T}$ being the event that the theory is false.

Interpretation: $P(T)$ is the *prior* probability of the theory being true. Given the evidence of the data, this prior probability is updated to arrive at the *posterior* probability, $P(T \mid D)$. The quantity, $P(D \mid T) / P(D)$ is the *updating* factor.

Application
The seller prices the cars, some at \$2500 and some at \$1000. He will always sell a good car for \$2500. He prices some of the bad cars at \$2500 and some at \$1000: the probability of a bad car being priced at \$2500 is μ and of it being priced at \$1000 is 1-μ; half of his cars are bad cars.

The buyer will accept a car priced at \$2500 with probability q and reject such a car with probability 1-q; the buyer will always buy a car priced at \$1000. The buyer believes that any car priced at \$2500 is a bad car with probability β and a good car with probability 1-β.

What is the probability that a bad car is sold for \$2500?

For the buyer:

$$\beta = P(B \mid p = 2500) = \frac{P(p = 2500 \mid B)P(B)}{p = 2500} = \frac{0.5\mu}{0.5\mu + 0.5} \tag{9}$$

Note: $P(p = 2500) = P(p = 2500 \mid G)P(G) + P(p = 2500 \mid B)p(B) = 0.5 + 0.5\mu$

If the buyer rejects the car priced at \$2500, his payoff is zero; if he accepts the \$2500 car then his payoff is:
$(3200 - 2500 - 1700)\beta + (3200 - 2500 - 200)(1 - \beta) = -1000\beta + 500(1 - \beta)$. In equilibrium, his expected payoff from rejection or acceptance of a \$2500 car must be the same implying: $500(1 - \beta) - 1000\beta = 0 \Rightarrow \beta = \frac{1}{3}$

Using this value of β to solve for μ: $\frac{1}{3} = \frac{0.5\mu}{0.5 + 0.5\mu} \Rightarrow \mu = \frac{1}{2}$

For the seller: If he offers a car for \$1000, his payoff is \$1000. If he offers a car for \$2500, his payoff is: $2500q + 0(1 - q)$. In equilibrium, the two payoffs are the same:
$1000 = 2500q \Rightarrow q = \frac{2}{5}$

What is the probability that a bad car is sold for $2500?

$$\begin{aligned} P(\text{sold at }\$2500 \mid B) &= P(\text{offered at }\$2500 \cap \text{accepted at }\$2500|\text{B}) \\ &= \frac{P(\text{offered at }\$2500 \cap \text{accepted at }\$2500 \cap \text{B})}{P(B)} \\ &= \frac{P(\text{offered at }\$2500 \cap B)P(\text{accepted at }\$2500)}{P(B)} \\ &= \frac{(1/2)\mu q}{(1/2)} = \frac{1}{2} \times \frac{2}{5} = 0.2 \end{aligned} \tag{10}$$

So seller can shift 20% of his stock of bad cars at the higher price of $2500.

Asymmetric Information and Discrimination: Application to Mortgage Lending

Individuals, who belong to one of two groups *Black (B)* or *White (W)* are looking for a loan. The likelihood with which they will repay the loan is $\theta \in [0,1]$. Banks define θ^* as the minimum degree of creditworthiness: a loan is approved for only those applicants for whom $\theta \geq \theta^*$. Unfortunately, lenders cannot observe θ. What they can observe from each applicant is a *signal*, *s*, which is correlated with θ. This signal may be thought of as a summary of all the information a bank collects on an applicant including: his income, nature of job, past credit history. Given the strength of the signal from an applicant, the bank estimates his expected creditworthiness: $q(s), dq/ds > 0$. It follows the rule that a loan application is approved if, and only if, $q(s) \geq q^* = \theta^*$.

The bank is said to *discriminate* against Black applicants if it requires them to meet a more stringent standard than that it does White applicants: $s_B^* > s_W^*$. This implies that Blacks have to be more creditworthy than Whites if they are to qualify for a loan.

Banks may discriminate against Blacks because they dislike Blacks. Gary Becker, *The Economics of Discrimination*, calls this taste-based discrimination. Banks are prepared to accept lower profits by turning away more creditworthy Black customers in favour of less creditworthy White customers and this reduction in profits is the "price" they pay for bigotry.

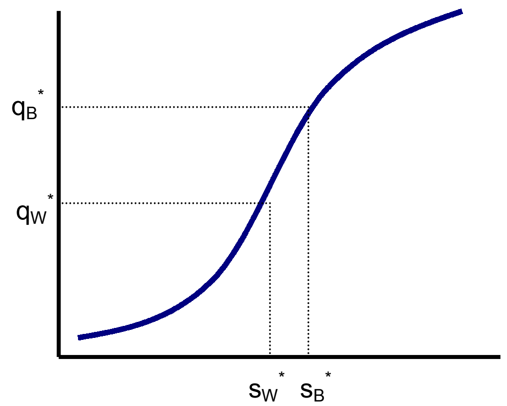

Figure 12
Taste-Based Discrimination (Bigotry)

In Figure 12, banks hold Black applicants to a more stringent underwriting standard ($s_B^* > s_W^*$); this implies that Black applicants have to cross a higher creditworthiness threshold ($q_B^* > q_W^*$).

Suppose now that the bank is not bigoted but it believes (or observes) that, *for the same signal strength, a Black applicant is less credit worthy than a White applicant.*

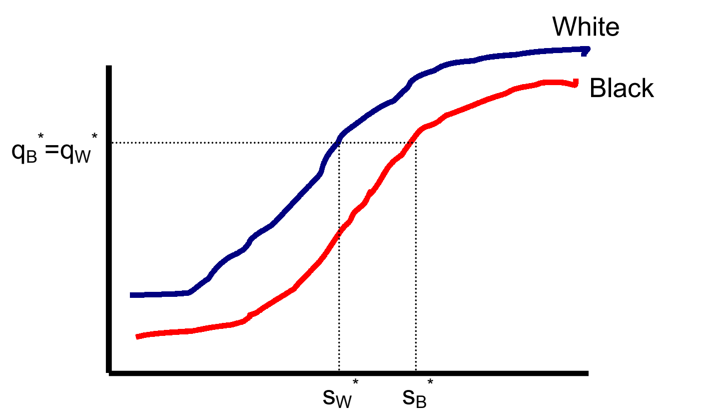

Figure 13
Statistical Discrimination

In Figure 13, the bank is practising *statistical discrimination*: Black and White applicants are set the *same* creditworthiness standards but the bank *discriminates* against Blacks by setting them a *higher* underwriting standard.

Now suppose we observe discrimination against Blacks in the sense that the compliance threshold for Blacks is higher than that for Whites: $s_B^* > s_W^*$. How can we tell whether this discrimination is due to “bigotry” or to “business necessity”. Under bigotry, with credit risk equally distributed amongst Blacks and Whites, we should observe a *lower* average default rate for Blacks; under “business necessity”, with average credit risk higher for Blacks than for Whites, we should observe the *same* average default rate for Blacks and Whites.