NBER WORKING PAPER SERIES

THE WELFARE ECONOMICS OF MORAL HAZARD

Richard Arnott

Joseph Stiglitz

Working Paper No. 3316

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 1990

This paper is part of NBER's research program in Taxation. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

NBER Working Paper #3316
April 1990

THE WELFARE ECONOMICS OF MORAL HAZARD

ABSTRACT

This paper shows that, except in certain limiting cases, competitive equilibrium with moral hazard is constrained inefficient. The first section compares the competitive equilibrium and the constrained social optimum in a fairly general model, and identifies types of market failure. Each of the subsequent sections focuses on a particular market failure.

Richard Arnott
Boston College
Carney 247
Chestnut Hill, MA 02167

Joseph Stiglitz
Stanford University
Department of Economics
Stanford, CA 94305

It is now widely recognized that the phenomenon of moral hazard, which arises whenever risk-averse individuals obtain insurance and their accident-avoidance activities cannot be perfectly monitored, is pervasive in the economy.[1] Since individuals do not bear the full consequences of their actions, incentives for accident avoidance tend to be less than if they did. This, in itself, does not imply that the market is (constrained) inefficient; to establish inefficiency, it needs to be shown that there is some intervention in the economy which would lead to a Pareto improvement. The object of this paper is to show that, in general, whenever moral hazard is present, market equilibrium is indeed "potentially" inefficient (i.e., assuming no costs of government intervention). The inefficiencies associated with market equilibrium with moral hazard take on a number of different forms, and this paper provides a taxonomy of these market failures.

Such a taxonomy is useful for several reasons. First, it helps in identifying different forms of government intervention which might yield Pareto improvements. Secondly, several of the inefficiencies which we identify in decentralized market economies can be thought of as externalities. There is a strong presumption that market economies respond to the existence of externalities by attempting to internalize them. Our analysis thus provides some insights into patterns of market structure which are otherwise difficult to explain, and, in the process of identifying the various forms of inefficiencies associated with moral hazard, enables us to ascertain the circumstances in which market "solutions" -- internalizing the informational externalities -- are more likely to be effective. Finally, the literature has identified certain limiting cases in which equilibrium is efficient. Our analysis helps to understand why these limiting cases are so special.

One of Arrow's great contributions was to show that the traditional competitive analysis, and hence the basic welfare theorems, can be extended to treat uncertainty, provided there is a complete set of insurance markets. The markets in which we are interested differ from Arrow-Debreu markets in an important way. Those markets provide insurance against <u>states of nature</u>, the occurrence of which is unaffected (by definition) by individuals' actions. Most insurance, however is for <u>events</u> (like hospitalization), the likelihood of which is affected by individuals' actions.[2]

Traditional results on the efficiency of market economies can be obtained even when the insured-against events are endogenous, so long as individuals' accident-prevention activities are observable. Inefficiencies arise only when neither the exogenous states of nature nor the individuals' accident-prevention activities are observable.

The fact that moral hazard alters the nature of market equilibrium has long been recognized. With moral hazard and complete insurance,[3] individuals have no incentive to avoid the accident; hence competitive markets typically entail incomplete insurance. But the existence of incomplete insurance does not imply that the market is necessarily Pareto inefficient, <u>given</u> the informational problems which are at the core of the moral hazard problem. One might argue (as Shavell [1979] and Pauly [1974], and several other authors have) that a competitive insurance firm takes into account how the level of insurance it provides affects the accident avoidance of those it insures and efficiently balances this against the benefits of risk reduction. This argument, as attractive as it has been to those who believe that markets must be efficient, is simply wrong. The expected profitability of an insurance contract depends on the actions taken by the insured, which are affected by his purchase of other commodities and other insurance (his savings, his income, etc.) all of which depend on prices. In an Arrow-Debreu economy, the externalities that such dependency on prices gives rise to do not cause inefficiency.[4] When moral hazard is present, however, we shall show that they do, and hence that competitive equilibrium is inefficient.

In earlier works, we showed that the nature of market equilibrium with moral hazard depends critically on whether the quantity of insurance which the individual purchases is or is not observable. In both cases, market equilibria are inefficient, but the nature of the inefficiencies and the potential role of the government differ. This paper focuses on the case where the quantity of insurance is observable.[5] In that case, we know that the equilibrium is characterized by a single firm providing the individual all of his insurance for a particular accident (see Arnott and Stiglitz [1987]).[6]

The basic source of the externalities is that if individuals take more care, some of the benefits from reduced accidents accrue to the insurance firm, not to the individual. If only one individual takes more care, his reduced accident probability will not be reflected in his premium; when all individuals take more care, they will be. Hence, policy interventions which increase care for all individuals -- provided they do not cost too much -- are desirable.

We can, accordingly, classify market failures by the avenue through which accident-avoidance activities are influenced. Accident avoidance may be affected by purchases of goods, which in turn are affected by prices. Firms, in setting their prices, fail to take this into account. This we call _the pecuniary externalities market failure_.

Accident avoidance also depends on individuals' income in different states. And individuals' income in different states depends on the insurance provided by different firms, including insurance provided against seemingly unrelated risks. Thus, the insurance which one firm provides may affect the profitability of insurance contracts offered by other firms against other risks. Associated with this is the _seemingly unrelated events market failure_.

Accident avoidance at date t is affected by wealth (and other state variables such as health) at date t; and this is affected by insurance provided in both prior and subsequent years. This is true even if accidents at different dates are independent events. Here a _wealth_ or _income effect_ gives rise to a seemingly unrelated events market failure.

One of the central implications of the efficiency of market economies is that cross-subsidies are not required. The fact that the amount of insurance purchased for one (seemingly unrelated) accident affects care in another suggests that there may be instances where it pays for one insurance policy to be taxed, to cross-subsidize another. This turns out in fact to be the case. That cross subsidies are not provided in competitive equilibrium causes the cross-subsidization market failure.

We establish these market failures using the following general approach. We first set up the optimization problem of the planner attempting to attain Pareto efficient outcomes and compare these with the market equilibrium. The planner, it turns out, takes into account certain terms which the market ignores, and while the market faces zero profit constraints for each firm, the planner has only an aggregate feasibility constraint. The shadow prices on goods derived in the planning problem turn out to differ from market prices, which implies the (potential) desirability of commodity taxation. And since the shadow prices on profits of different firms will typically differ, cross subsidies are in general desirable.

The general approach allows us to identify the limiting cases where market equilibrium is Pareto efficient. The expressions characterizing market equilibrium and Pareto efficiency turn out to be identical when certain behavioral responses are absent, i.e. certain derivatives are zero. By focusing on special cases of our general model, we obtain insights into the highly unusual conditions under which these behavioral responses do not appear. Unfortunately, much of the earlier literature, attempting to establish the efficiency of market economies, focused precisely on those special cases, for instance, where there is a single consumption good and a single accident, and where each individual purchases all of his insurance from a single competitive insurer.[7]

The paper is divided into four sections. The next section sets up the general model, while the following sections examine in greater detail several of the market failures we have identified.

1. A Fairly General Model

There is a group of individuals identical in all respects[8], faced with the possibility of several mutually exclusive outcomes or events,[9] $i \epsilon I$. Each event may affect an individual's welfare directly and/or result in the loss of a given bundle of consumer goods, d_i. It is assumed that an individual receives utility from one vector of consumer goods that he purchases prior to the outcome, $\bar{c}$, and another vector of consumer goods that he purchases subsequent to the outcome, $\hat{c}_i$, and disutility from various types of accident-prevention effort, e; i.e., utility with outcome i is given by

$$u_i = u_i(\bar{c},\ \hat{c}_i - d_i,\ e) \ . \tag{1.1}$$

The probability of each outcome depends on both $\bar{c}$ and e,

$$p_i = p_i(\bar{c}, e) \ . \tag{1.2}$$

The individual is assumed to supply a fixed quantity of labor, as well as a fixed number of units of non-labor factors of production. The economy is large, and different individuals' outcome probabilities are independent. Thus, there is no aggregate uncertainty. Furthermore, it is assumed that there are constant returns to scale in production, and that production is competitively organized. As a result, with labor income as numéraire, the return to non-labor factors of production or producer profits, $\hat{\pi}$, can be expressed as a function of producer prices, r , i.e.

$$\hat{\pi} = \hat{\pi}(r) \ . \tag{1.3}$$

The individual's before-insurance income is therefore

$$x = l + \hat{\pi}(r) \ . \tag{1.4}$$

It is assumed that insurers can observe neither an individual's effort nor his total purchase of various commodities. This is the source of moral hazard. As a result, insurance policies cannot be based on these magnitudes. To simplify the analysis, we assume that there is a separate insurance policy for each outcome. For a subset of these outcomes, $i \in I'$, we assume that insurance purchases are observable, in which case (see Arnott and Stiglitz [1987]) exclusive insurance contracts specify both the quantity of insurance the individual is to purchase and its price; equivalently, policy i is characterized by β_i , the premium, and γ_i , the gross (of the contract's own premium) payout. For the remaining outcomes, $i \in I/I'$, an individual's total insurance purchases are not observable, and we assume that the individual can purchase as much insurance as he wants at the price (premium/gross payout ratio) θ_i. Thus, where z_j is the net insurance payout with outcome j,

$$z_j = \gamma_j - \sum_{i \in I'} \beta_i - \sum_{i \in I/I'} \gamma_i \theta_i \ . \tag{1.5}$$

Also, the expected profits of the i^{th} insurance policy are

$$\pi_i = \begin{cases} \beta_i - \gamma_i p_i & \text{for } i \epsilon I' \\ \gamma_i \theta_i - \gamma_i p_i & \text{for } i \epsilon I/I' \ . \end{cases} \tag{1.6}$$

Hence, the total expected profits of all policies together are

$$\pi = \sum_{i \epsilon I} \pi_i = \sum_{i \epsilon I'} \beta_i + \sum_{i \epsilon I/I'} \gamma_i \theta_i - \sum_{i \epsilon I} \gamma_i p_i = - \sum_{i \epsilon I} p_i z_i \ . \tag{1.7}$$

To be consistent with earlier assumptions about the composition of individuals' incomes and to keep the model general equilibrium in nature, these insurance profits accrue to the government.[10]

We assume that transactions, but not individuals' total purchases, of goods are observable. This allows linear, but not non-linear, commodity taxation. Where q denotes the vector of consumer prices, from (1.4) and (1.5), the individual's budget constraint with outcome j is

$$x + z_j = 1 + \hat{\pi}(r) + \gamma_j - \sum_{i \epsilon I'} \beta_i - \sum_{i \epsilon I/I'} \gamma_i \theta_i = \bar{q}\bar{c} + \hat{q}\hat{c}_j \ . \tag{1.8}$$

The individual maximizes expected utility subject to his outcome-contingent budget constraints, i.e.

$$\max_{e, \bar{c}, (\hat{c}_i)_{i \epsilon I}, (\gamma_i)_{i \epsilon I/I'}} EU = \sum_{i \epsilon I} p_i(\bar{c}, e) u_i(\bar{c}, \hat{c}_i - d_i, e) \tag{1.9}$$

subject to (1.8), the outcome-dependent budget constraints.

Thus, using obvious vector notation,

$$\bar{c} = \bar{c}(r,\gamma,\beta,\theta,q) , \qquad \hat{c}_i = \hat{c}_i(r,\gamma,\beta,\theta,q) \quad \text{for} \quad i \epsilon I$$

$$e = e(r,\gamma,\beta,\theta,q) , \qquad EU = V(r,\gamma,\beta,\theta,q) , \text{ and} \tag{1.10}$$

$$\gamma_i = \gamma_i(r,\gamma,\beta,\theta,q) \quad \text{for} \quad i \epsilon I/I' ,$$

where γ denotes the vector of γ_i <u>for $i \epsilon I'$</u>. If β, θ, q, and γ are fixed, then r is determined by market-clearing in the goods, factor, and insurance markets. Hence we may write

$$r = r(\gamma,\beta,\theta,q) . \tag{1.11}$$

Total expected tax revenue is

$$T = (\bar{q}-\bar{r})\bar{c} + \sum_{i \epsilon I} p_i(\hat{q}-\hat{r})\hat{c}_i . \tag{1.12}$$

The feasibility constraint for the economy is that

$$\pi + T \geq 0 . \tag{1.13}$$

Also, using (1.7), (1.10), (1.11), and (1.12), we have

$$EU = V(r(\gamma,\beta,\theta,q),\gamma,\beta,\theta,q), \tag{1.14a}$$

$$\pi_i = \pi_i(r(\gamma,\beta,\theta,q),\gamma,\beta,\theta,q) , \text{ and} \tag{1.14b}$$

$$T = T(r(\gamma,\beta,\theta,q),\gamma,\beta,\theta,q) . \tag{1.14c}$$

If the planner has $\eta = (\gamma,\beta,\theta,q)$ as instruments,[11] the social optimum is characterized by

$$\max_{\eta} \mathcal{L} = V(r(\eta),\eta) + \lambda[\sum_{i\epsilon I} \pi_i(r(\eta),\eta) + T(r(\eta),\eta)] \ , \tag{1.15}$$

the first-order conditions of which are[12]

$$\underset{}{\frac{\partial V}{\partial \eta_j}} + \lambda \left[\underset{(a)}{\sum_{i\epsilon I} \frac{\partial \pi_i}{\partial \eta_j}} + \underset{(b)}{\frac{\partial T}{\partial \eta_j}} \right] + \sum_{\ell \epsilon L} \left[\underset{(c_1)}{\frac{\partial V}{\partial r_\ell}} + \lambda \left(\underset{(c_2)}{\sum_{i\epsilon I} \frac{\partial \pi_i}{\partial r_\ell}} + \underset{(c_3)}{\frac{\partial T}{\partial r_\ell}} \right) \right] \frac{\partial r_\ell}{\partial \eta_j} = 0 \tag{1.16}$$

where L is the set of all commodities (i.e. includes both $\bar{c}$ and $\hat{c}_i$ commodities).

This needs to be contrasted with the market equilibrium (assuming it exists), in which each competitive insurance firm can be viewed as providing a single policy, choosing the parameters of the policy to maximize the individual's expected utility subject to its at least breaking even, and taking the vector of producer prices and the parameters of all other insurance contracts as given. The competitive equilibrium can be shown to be the solution to the following set of equations:

$$q = r(\gamma,\beta,\theta,q) \ , \tag{1.17}$$

i.e. consumer prices equal producer prices; in the normal case (see Arnott and Stiglitz [1987] for a discussion of the other cases)

$$\frac{\partial V}{\partial \gamma_i} = 0 \tag{1.18}$$

for price insurance contracts (i.e. $i\epsilon I/I'$) , which can be shown to imply that the private marginal utility of income must be the same for all outcomes covered by price insurance contracts; and

$$\frac{\partial V}{\partial \gamma_j} + \lambda_j \frac{\partial \pi_j}{\partial \gamma_j} = 0 \qquad \frac{\partial V}{\partial \beta_j} + \lambda_j \frac{\partial \pi_j}{\partial \beta_j} = 0 \tag{1.19}$$

for quantity-constrained insurance policies, where λ_j is the Lagrange multiplier associated with the j^{th} insurance firm's profit constraint.

Comparing (1.16) with (1.17) - (1.19), we can see the major sources of market failure noted in the introduction. The market ignores: (a) the effect of a change in the terms of the policy (or a price) on profits of other policies; (b) the effect of the change on tax revenues; (c) the effect of the change on producer prices, and the consequent effect on utility, either directly, or indirectly through the effect on producer and insurance profits and on tax revenues. We now describe each of these market failures more fully.

a. <u>The insurance externality with quantity rationing</u>

In general, the quantity of insurance purchased for the j^{th} outcome affects the actions individuals undertake, which affect the likelihood of the occurrence of other events and hence the insurance profits for all other outcomes (as well as profits in other industries). The competitive firm ignores this, while the social planner takes it into account.

Formally, this can be seen in the comparison between (1.19) and (1.16). We can rewrite (1.19) to read

$$\frac{\partial V/\partial \gamma_j}{\partial V/\partial \beta_j} = \frac{\partial \pi_j/\partial \gamma_j}{\partial \pi_j/\partial \beta_j}, \quad j \epsilon I'$$

and (1.16), to read

$$\frac{\partial V/\partial \gamma_j}{\partial V/\partial \beta_j} = \frac{\partial \pi_j/\partial \gamma_j + \sum_{i \neq j} \frac{\partial \pi_i}{\partial \gamma_j} + \frac{\partial T}{\partial \gamma_j} + \sum_{\ell \epsilon L} \frac{\partial r_\ell}{\partial \gamma_j}\left[\left(\sum_{i \epsilon I} \frac{\partial \pi_i}{\partial r_\ell} + \frac{\partial T}{\partial r_\ell}\right) + \frac{1}{\lambda}\frac{\partial V}{\partial r_\ell}\right]}{\partial \pi_j/\partial \beta_j + \sum_{i \neq j} \frac{\partial \pi_i}{\partial \beta_j} + \frac{\partial T}{\partial \beta_j} + \sum_{\ell \epsilon L} \frac{\partial r_\ell}{\partial \beta_j}\left[\left(\sum_{i \epsilon I} \frac{\partial \pi_i}{\partial r_\ell} + \frac{\partial T}{\partial r_\ell}\right) + \frac{1}{\lambda}\frac{\partial V}{\partial r_\ell}\right]}, \quad j \epsilon I'.$$

The effect with which we are concerned here arises when r is constant and $T = 0$. Then (1.16) becomes

$$\frac{\partial V/\partial \gamma_j}{\partial V/\partial \beta_j} = \frac{\partial \pi_j/\partial \gamma_j + \sum_{i \neq j} (\partial \pi_i/\partial \gamma_j)}{\partial \pi_j/\partial \beta_j + \sum_{i \neq j} (\partial \pi_i/\partial \beta_j)}, \quad j \epsilon I'.$$

The distortion arises from the effect of insurance on other insurance firms' profits, which is captured by the terms

$$\sum_{\substack{i \neq j \\ j \epsilon I'}} \frac{\partial \pi_i}{\partial \gamma_j} \quad \text{and} \quad \sum_{\substack{i \neq j \\ j \epsilon I'}} \frac{\partial \pi_i}{\partial \beta_j}.$$ [13]

To emphasize the fact that this arises even when the events which are being insured are apparently unrelated, we refer to this as the "seemingly unrelated events market failure."

b. The pecuniary externalities market failure

(1.16) also suggests that it may in general be desirable to impose a tax on the commodities purchased. To see this, consider the case in which producer prices are constant. Then, from (1.16), in the absence of commodity taxes,

$$\left.\frac{\partial \mathcal{L}}{\partial q_\ell}\right|_{q=r} = \frac{\partial V}{\partial q_\ell} + \lambda \left[\sum_{i \epsilon I} \frac{\partial \pi_i}{\partial q_\ell} + \frac{\partial T}{\partial q_\ell} \right] .$$

From the individual's maximization problem, (1.9),

$$\frac{\partial V}{\partial q_\ell} = - \sum_{i \epsilon I} \epsilon_i p_i c_{i\ell} ,$$

where ϵ_i is the individual's marginal utility of income under event i .

Since $\left.\frac{\partial T}{\partial q_\ell}\right|_{q=r} = \sum_{i \epsilon I} p_i c_{i\ell}$ from (1.12), then,

$$\left.\frac{\partial \mathcal{L}}{\partial q_\ell}\right|_{q=r} = - \sum_{i \epsilon I} (\epsilon_i - \lambda) p_i c_{i\ell} + \lambda \sum_{i \epsilon I} \frac{\partial \pi_i}{\partial q_\ell} .$$

Thus, if $\lambda = \epsilon_i$ for all i and $\sum_{i \epsilon I} \frac{\partial \pi_i}{\partial q_\ell} = 0$, it is not optimal to tax good ℓ . But typically $\lambda \neq \epsilon_i$ for all i , and from (1.6)

$$\frac{\partial \pi_i}{\partial q_\ell} = \begin{cases} -\gamma_i \dfrac{\partial p_i}{\partial q_\ell} & \text{for } i \epsilon I' \\ -\gamma_i \dfrac{\partial p_i}{\partial q_\ell} + (\theta_i - p_i) \dfrac{\partial \gamma_i}{\partial q_\ell} & \text{for } i \epsilon I/I' \end{cases}$$

which is not in general equal to zero. Thus, it appears generally desirable to tax or subsidize commodities which affect accident probabilities, because of their direct and indirect effects on accident-avoidance activities and the consumption of commodities which directly affect the likelihood of an accident. That this is indeed the case is demonstrated in Arnott and Stiglitz [1986].

c. The income effect

In the previous subsection, producer prices, and hence incomes, were held constant. In general, a change in γ, β , θ , and q will affect producer prices and hence profits. These changes in profits will affect care, and hence will have further repercussions on profits in the insurance industry. This effect, captured in (1.16) by terms in $\partial r_\ell/\partial \eta_j$, is obviously neglected by the insurance firm in choosing its profit-maximizing contract, and is a second source of pecuniary externalities.

d. The insurance externality without quantity rationing

Though the government cannot directly control the quantity of insurance that the individual purchases against those outcomes for which individual total insurance coverage is unobservable, it can affect the quantity purchased by taxing such insurance. This will reduce the inefficiencies stemming from the inadequate precautions taken as a result of excessive insurance purchased against those outcomes. To isolate this effect, set T = 0 and take producer prices as fixed. From the consumer's maximization problem for $j \epsilon I/I'$, $\frac{\partial V}{\partial \theta_j} = -\sum_{i \epsilon I} p_i \epsilon_i \gamma_j$, while from (1.7),

$$\frac{\partial \pi_j}{\partial \theta_j} = \gamma_j - \sum_{i \epsilon I} \gamma_i \frac{\partial p_i}{\partial \theta_j} + \sum_{i \epsilon I/I'} \frac{\partial \gamma_i}{\partial \theta_j} (\theta_i - p_i).$$ From (1.16), at the

competitive equilibrium

$$\left. \frac{\partial \mathcal{L}}{\partial \theta_j} \right|_{\theta = p} = \frac{\partial V}{\partial \theta_j} + \lambda \frac{\partial \pi_j}{\partial \theta_j}$$

$$= (\lambda - \sum_{i \epsilon I} p_i \epsilon_i) \gamma_j - \lambda \sum_{i \epsilon I} \gamma_i \frac{\partial p_i}{\partial \theta_j}$$

which will not in general equal zero.

e. No cross-subsidization

To isolate this market failure, ignore taxes. The market equilibrium generates non-negative profits for the insurance policies covering each outcome. If the government were constrained to at least break even in each policy rather than on all policies together, (1.15) would become

$\max_{\eta} \mathcal{L} = V(r(\eta),\eta) + \sum_{i \epsilon I} \lambda_i \pi_i (r(\eta),\eta)$, and since $\lambda_j \neq \lambda_k$ in general, a Pareto improvement could be made by transferring funds from one insurance policy to another.

The careful reader at this point may be wondering: Some of these effects arise in the classic competitive economy, but do not give rise to inefficiency there. The actions taken by one firm may affect producer prices, and hence the profits faced by other firms, but inefficiency does not result; and firms face separate budget constraints, but there is no cross-subsidization market failure. In the remaining sections, we explain why these effects give rise to market failure when moral hazard is present, even though they do not when it is absent, and provide some insight into the nature and direction of the biases generated by moral-hazard-induced inefficiencies. Our discussion below focuses on three of the market failures, the seemingly unrelated events market failure, the cross-subsidization market failure, and pecuniary externalities market failures arising from income effects. In three companion papers [Arnott and Stiglitz, 1986, 1988c, 1988d], we describe in greater detail the other market failures.

2. Seemingly Unrelated Events Market Failure

In a previous paper (Arnott and Stiglitz [1987]) in which there was a single fixed-damage accident and insurance policies were observable, we showed that competitive equilibrium entails exclusive contracts--each individual purchases all his insurance from a single firm (see Footnote 6). Since there are many firms offering such policies, exclusivity is consistent with competitiveness. Here we extend that result to show that when there are several risks and insurance policies are observable, an individual should have all his insurance needs served by a single firm, even when the risks are seemingly unrelated; we term this "extended exclusivity".

In the previous section, we argued that the actions of a firm offering insurance against one type of accident generally affect the profits of firms offering insurance against other types of accident. Exclusivity is required to internalize these externalities. We did not, however, prove that each firm's actions affect the profitability of other firms, nor did we investigate the direction of the biases introduced by these inefficiencies. We do so in this section.

That there are important interactions among different risks seems clear: An individual who drinks excessively will have a higher likelihood of an automobile accident and a higher likelihood of hospitalization. Increasing the degree of hospitalization insurance means that the magnitude of an individual's total losses from drinking will be reduced, and he may therefore be induced to drink more, increasing the automobile accident rate and lowering profits of automobile insurance. In this case, there would be a negative externality--the market would be characterized by excessive hospitalization insurance (and by symmetry) excessive automobile insurance.

What is not so apparent, however, is that even when there is apparently no relation between the events (and the accident-avoidance activities), there is an interdependence. Not surprisingly, the effects are complicated, and it is not an easy task to sign them. Though the calculations are complex, the analytic approach we take is simple: We ascertain the effect of a change in insurance against one accident on the likelihood of the occurrence of other accidents; that is, we calculate the derivative of the probability of accident j with respect to the terms of insurance against accident i; so long as this derivative is not zero, the market equilibrium is inefficient.

To simplify the discussion, we consider a world in which there are only two types of fixed-damage accident, 1 and 2. There are four possible outcomes: no accident, accident 1, accident 2, and both accidents. Suppose that a representative competitive insurance firm 1 provides insurance against accident 1, and a representative competitive insurance firm 2 against accident 2. To simplify further, we assume that there is a single commodity, the consumption of which has no effect on accident probabilities. The probabilities of the accidents depend only on the levels of various types of accident-prevention effort.

There are three avenues through which a budget-balancing perturbation of firm 1's contract can affect the profitability of firm 2. First, one or more of the types of accident-prevention effort may directly affect (i.e., enter the probability-of-accident functions for) both accidents.[14] Suppose firm 1 offers more insurance. This will affect the effort levels chosen by the individual, which will affect the probability of accident 2 and hence the profitability of firm 2. Second, even when the various types of accident-prevention effort are accident-specific, the level of one type of effort may affect the marginal disutility of other types of effort[15] and hence the levels of these other types of effort that the individual chooses. When firm 1 offers more insurance, this may affect the marginal disutility of a type of effort that influences the profitability of firm 2.[16] Third, even with neither of these effects operative, an increase in the provision of insurance against one accident may make the individual more or less complacent and careless in preventing the other, because it affects the differences in average marginal utilities between those states where the accident does and does not occur. The nature of this effect depends critically on the correlations between the two accidents, as we shall shortly see.

The first two avenues, discussed above, through which the contract offered by firm 1 can affect the profitability of firm 2, are obvious. We shall therefore model only the last. Let $j=1,2$ index both the firm and accident type, and i index the outcome with: $i = 0$, no accident; $i = 1$, accident 1 ; $i = 2$, accident 2 ; and $i = 3$, both accidents. Corresponding to each accident, there is a single unobservable, accident-specific type of accident-prevention effort e_j .

To simplify the algebra, we assume that there is complete symmetry between the two accidents. Let $P_j = P(e_j)$ be the probability of accident j, with $P' < 0$, $P'' > 0$, and p_i be the probability of outcome i, which are related as follows

$$p_0 = 1 - P_1 - P_2 + \Omega P_1 P_2 \qquad p_1 = P_1(1-\Omega P_2)$$
$$p_2 = P_2(1-\Omega P_1) \qquad p_3 = \Omega P_1 P_2 . \tag{2.1}$$

The parameter Ω captures the correlation between the accidents; $\Omega = 0$ corresponds to mutually exclusive accidents, $\Omega = 1$ to statistically independent accidents, and $\Omega > 1$ to positively-correlated accidents.

The expected utility function has the special form

$$EU = \sum_{i=0}^{3} u(y_i)p_i - \sum_{j=1}^{2} e_j , \tag{2.2}$$

where y_i is net-of-insurance income (consumption) with outcome i. We term this a separable, event-independent expected utility function--event-independent since the accidents do not affect the utility function directly.

Note that since the effort types are accident-specific, the first avenue of interdependence is excluded. Since the marginal disutility of each effort type is independent of the level of the other effort type, the second avenue of interdependence is excluded too.

Firm 1 and firm 2 both offer quantity-constrained contracts of the form (α_j, β_j), where α_j is the net (of premium) insurance payout if accident j occurs, and β_j is the insurance premium for accident j. Thus,

$$y_0 = w - \beta_1 - \beta_2 \qquad y_1 = w - d + \alpha_1 - \beta_2$$
$$y_2 = w - \beta_1 - d + \alpha_2 \quad y_3 = w - d + \alpha_1 - d + \alpha_2 \,, \tag{2.3}$$

where w is the individual's pre-insurance income when an accident does not occur, and d is the size of the fixed-damage accident.

The first-order conditions of the individual's effort choice decision are

$$[(1-\Omega P_j)(u_{j'} - u_0) + \Omega P_j(u_3 - u_j)]P'_{j'} = 1 \quad j, j' = 1,2;\ j \neq j' \,. \tag{2.4}$$

We assume an interior solution, which may be written as

$$e_j = e_j(\alpha_1, \alpha_2, \beta_1, \beta_2) \quad j = 1,2. \tag{2.5}$$

Expected utility (substituting (2.5) into (2.2)) is

$$EU = V(\alpha_1, \alpha_2, \beta_1, \beta_2) \,. \tag{2.6}$$

Firm j's problem is to

$$\max_{(\alpha_j, \beta_j)} V(\alpha_1, \alpha_2, \beta_1, \beta_2) \quad \text{s.t. } B_j = \beta_j(1-P_j) - \alpha_j P_j \geq 0 \,,\ j = 1,2 \,. \tag{2.7}$$

Firm j must at least break even on policy j, and takes $\alpha_{j'}$ and $\beta_{j'}$ as given. Where λ_j is the multiplier on B_j, the first-order conditions are

$$\frac{\partial V}{\partial \alpha_j} + \lambda_j \frac{\partial B_j}{\partial \alpha_j} = 0 \qquad \frac{\partial V}{\partial \beta_j} + \lambda_j \frac{\partial B_j}{\partial \beta_j} = 0 . \tag{2.8}$$

To ascertain whether competitive equilibrium entails over- or under-provision of insurance, we perform the following exercise: We ask if a budget-balancing increase in insurance offered by firm 1 in the neighborhood of the competitive equilibrium, holding firm 2's contract fixed, stimulates or discourages the effort expended by the individual to prevent accident 2. If the increase stimulates e_2, then the social benefit from the increase exceeds the private benefit, and it is desirable that firm 1 offer more insurance; hence, competitive equilibrium entails under-insurance. Let $\left.\frac{de_2}{d\alpha_1}\right|_{B_1}$ denote the change induced from an increase in α_1, holding α_2 and β_2 constant, but allowing β_1 to change to maintain budget balance. Then

$$\left.\frac{de_2}{d\alpha_1}\right|_{B_1} = \frac{\partial e_2}{\partial \alpha_1} + \frac{\partial e_2}{\partial \beta_1} \left.\frac{d\beta_1}{d\alpha_1}\right|_{B_1} \begin{matrix} > \\ < \end{matrix} 0 \Rightarrow \begin{matrix} \text{under-insurance} \\ \text{over-insurance} \end{matrix} . \tag{2.9}$$

Now, at the competitive equilibrium, from (2.8), $\left.\frac{d\beta_1}{d\alpha_1}\right|_{B_1} = \left.\frac{d\beta_1}{d\alpha_1}\right|_{V}$.

Hence,

$$\left.\frac{de_2}{d\alpha_1}\right|_{B_1} = \frac{\partial e_2}{\partial \alpha_1} + \frac{\partial e_2}{\partial \beta_1}\left.\frac{d\beta_1}{d\alpha_1}\right|_V . \tag{2.10}$$

$\frac{\partial e_2}{\partial \alpha_1}$ and $\frac{\partial e_2}{\partial \beta_1}$ are obtained from total differentiation of (2.4), while $\left.\frac{d\beta_1}{d\alpha_1}\right|_V$ is obtained from (2.2). Where $P = P_1 = P_2$ and $u. = u_1 = u_2$ since the two accidents are symmetric, substitution of these results gives

$$\left.\frac{de_2}{d\alpha_1}\right|_{B_1} = \frac{1}{\Delta}\left\{\begin{vmatrix} -\Omega PP'(u_3'-u_.') & \Omega(P')^2(u_3-2u_.+u_0) \\ -((1-\Omega P)u_.'+\Omega Pu_3')P' & \frac{P''}{P'} \end{vmatrix} + \begin{vmatrix} (1-\Omega P)P'(u_.'-u_0') & \Omega(P')^2(u_3-2u_.+u_0) \\ -((1-\Omega P)u_0'+\Omega Pu_.')P' & \frac{P''}{P'} \end{vmatrix} \left(\frac{P(1-\Omega P)u_.'+\Omega P^2u_3'}{(1-2P+\Omega P^2)u_0'+P(1-\Omega P)u_.'}\right)\right\}, \tag{2.11}$$

where $\Delta > 0$ from the second-order conditions of the individual's effort-choice problem.

We shall consider two special cases.

a. Mutually exclusive accidents ($\Omega=0$)

We obtain

$$\left.\frac{de_2}{d\alpha_1}\right|_{B_1} = (u_.'-u_0')\frac{P''}{\Delta}\left(\frac{Pu_.'}{(1-2P)u_0'+Pu_.'}\right) > 0 , \tag{2.12}$$

since $u_.' - u_0' > 0$ (with moral hazard, insurance is only partial, which implies that $y_0 > y_2$) and $P'' > 0$. Thus, with mutually exclusive accidents, the competitive equilibrium entails underinsurance. The reason for this is as follows: With mutually exclusive accidents, $p_0 = 1 - P_1 - P_2$, $p_1 = P_1$, $p_2 = P_2$, and $p_3 = 0$, and from (2.4), the first-order condition for e_2 is

$$-(u_0 - u_2)\, P_2' - 1 = 0 .$$

In deciding on e_2, the individual will compare the utility in the no-accident event with that in the accident 2 event (since e_2 does not affect the probability of accident 1). The increase in insurance against accident 1 decreases y_0 and y_2 by the same amount (recall (2.3)). Because of diminishing marginal utility of income, and since $y_0 > y_2$, this increases $u_0 - u_2$, which stimulates effort.

b. <u>Statistically independent accidents</u> ($\Omega=1$)

We obtain

$$\left.\frac{de_2}{d\alpha_1}\right|_{B_1} = \frac{1}{\Delta}\left\{ \frac{PP''}{(1-P)u_0' + Pu_\cdot'}\left((u_\cdot')^2 - u_0'u_3'\right) + \frac{(P')^3}{1-P}\left((1-P)u_\cdot' + Pu_3'\right)(u_3 - 2u_\cdot + u_0) \right\} . \tag{2.13}$$

The second term in the curly brackets is unambiguously positive (since $(P')^3 < 0$ and $u_3 - 2u_\cdot + u_0 < 0$ (concavity of u)), while the first term is positive if absolute risk aversion is constant or increasing over the relevant range $((u_\cdot')^2 - u_0'u_3' > 0 \Longleftrightarrow 2 \ln u_\cdot' \geq \ln u_0' + \ln u_3' \Longleftrightarrow \ln u'$ is concave in y (since $y_2 = y_1 = \frac{y_0 + y_3}{2}$) $\Longleftrightarrow$ constant or increasing absolute risk aversion). Hence, there is a presumption that the market also under-provides insurance in the case of statistically independent accidents. If, however, absolute risk aversion is sufficiently decreasing, the market will over-provide insurance.[17]

The general point is that an increase in the amount of insurance provided by a firm against one type of accident will, by altering the marginal utilities of income of several outcomes, affect the marginal benefit of effort, and hence the effort level chosen, in preventing other accidents. As a result, insurance firms generate external effects that can be internalized if each individual purchases all his insurance from one agent--the extended exclusivity requirement for efficiency. If individuals purchase different types of insurance from different firms, the market will be inefficient, but the direction of bias--what types of accident the market will over-insure and what types it will under-insure--is in general ambiguous, depending in a complex way on such factors as the correlation between the accidents, the risk-aversion properties of the utility function, and the characteristics of the accident-prevention technology.

c. Dynamic interactions

Thus far, we have discussed different types of risk at a point in time. The rule that exclusivity is a necessary condition for efficiency applies as well to the provision of insurance against the same type of risk and different types of risk over time. Thus, not only should an individual's insurance needs be served by a single insurance agent at a point in time, but the individual should also have the same insurance agent through time. This is an important implication of our analysis. For most types of risks, the probability of accident depends on the value of some imperfectly observable[18] stock or state variables--weight, state of health, education, savings--over which the individual has at least partial control and which are influenced by the amount of insurance provided through time. An agent who provides insurance over only part of the insured's life will neglect the effect of the insurance he provides on the value of these stock variables, and hence on accident probabilities, before and after the period in which he is the insurer.

Traditional theories have stressed the importance of precautionary savings. Since these savings are not earmarked, they serve to "insure" simultaneously against all risks. The provision of greater market insurance against some risk (say fire) reduces the need for savings; but at reduced level of savings, individuals will tend to take greater care (at fixed levels of insurance) against other risks. With this effect, an increase in insurance against one risk will normally have a positive effect on the profitability of insurance against other risks, which implies that precautionary savings lead to under-insurance.

An extended example of these dynamic externalities is provided by Arnott and Stiglitz [1985].

d. The seemingly unrelated events market failure.

Since we have assumed that there are no administrative costs in the provision of insurance and that individuals are perfectly informed concerning the menu of contracts being offered, competition should result in extended exclusivity--a firm which offers the socially optimal exclusive contract covering all of one individual's insurance needs throughout his life would drive all other firms out of business. Assume, more realistically, that insurance administrative costs are characterized by diseconomies of scope, so that competition will result in each firm specializing in the provision of only a subset of types of insurance. Moreover, for a variety of reasons, intertemporal exclusivity is unlikely; individuals would have to sign up at birth with an insurance firm that would cover them for all risks (even those which, at the time, they are not fully aware of) throughout their lives, regardless of where they subsequently choose to live or the occupation or lifestyle they decide to pursue. In either of these cases, the externalities we have identified in this section would be present. Not only would the market over- or under-provide insurance, but also insurance firms would in general be over- or under-specialized. Thus, we have identified a genuine potential market failure. The inefficiency arises when accidents are related, but since its appearance is most surprising when the accidents appear unrelated, we term it the <u>seemingly unrelated events market failure</u>.

In fact, individuals do obtain insurance from a variety of sources; typically, one obtains market insurance from more than one carrier; and if one is sick, one usually gets compensated sick leave from work and medical care that is subsidized by the government.

Whether insurance markets are well-described by a competitive model such as the one we have presented is moot, but whatever the market structure, in the absence of extended exclusivity, the seemingly unrelated events market failure is present.

3. <u>The Cross-Subsidization Inefficiency</u>

A central property of competitive models is that no cross-subsidies between firms are needed. Consider two constant returns industries producing consumer goods. Assume that an infinitesimal specific tax is imposed on consumer good i, and that the revenues are used to finance a subsidy on good j. Good j's price is lowered, and i's price is raised; the gain on one account is just equal to the loss on the other. There are, of course, further repercussions: consumption of other goods, supply of labor, etc. will all change. But since all individuals are maximizing their utility, these adjustments have second-order welfare effects (because of the envelope theorem), and any consequent price changes have, at most, redistributive effects.[19] Any subsidies/taxes that are not infinitesimal have further distortionary effects, and are welfare-reducing.

In contrast, with moral hazard, cross-subsidization is in general desirable. The basic idea is that the transfer of a dollar from one firm to another will alter the general equilibrium of the economy, including individuals' effort levels and hence the "deadweight loss" associated with moral hazard (i.e. the loss relative to the equilibrium in which effort is observable). The tax on one insurance policy leads to a price increase, a decrease in the quantity of insurance purchased, and an increase in accident-avoidance effort against that risk. The subsidy on the other insurance policy has qualitatively the opposite effects. These effort effects will not in general be offsetting. There may, of course, be other general equilibrium effects, for instance on the prices of various commodities and the levels of consumption, but (apart from any further induced effects on effort) these have second-order welfare consequences, because of the fact that individuals are utility-maximizing and firms are profit-maximizing.

We shall demonstrate the desirability of cross-subsidization for insurance firms, but it holds generally. This gives a further argument for exclusivity, since if an individual purchases all his insurance from a single carrier, the carrier can cross-subsidize between contracts. When insurance against different risks is provided by different carriers, due to regulation, diseconomies of scope, etc., then cross-subsidization between carriers is potentially welfare-improving. This is true even when the government cannot impose taxes or subsidies directly on the quantity of insurance purchased by an individual, because of costly monitoring, for example.

To establish the desirability of cross-subsidization, we employ a model in which there are two possible accidents, 1 and 2. The individual commits himself to undertaking either project 1 or project 2 with equal probability (e.g., project 1 could be a sunny-day project, and project 2 a rainy-day project) before expending any accident-prevention effort. Accident 1 can occur only if the individual embarks on project 1, while accident 2 only with project 2. Thus, the two accidents are mutually exclusive. But here they are *ex ante* exclusive--the roll of the die to determine which of the accidents cannot occur is made *before* the individual makes his effort decision; while in the previous section, the mutual exclusivity was *ex post*. We denote variables associated with project 1 by a $\hat{}$ and with project 2 with a $\tilde{}$. Thus, the expected utility function is

$$EU = \frac{1}{2}\{\hat{u}_0(1-\hat{P})+\hat{u}_1\hat{P}-\hat{e}\} + \frac{1}{2}\{\tilde{u}_0(1-\tilde{P})+\tilde{u}_1\tilde{P}-\tilde{e}\} \quad . \tag{3.1}$$

Subsequently, we shall drop the $\frac{1}{2}$'s . We impose a lump-sum tax S (possibly negative) on contracts providing insurance against accident 1, and use the proceeds to finance a lump-sum subsidy of S for contracts providing insurance against accident 2. Thus, contracts 1 and 2 have the budget constraints

$$\hat{\beta}(1-\hat{P}) - \hat{\alpha}\hat{P} - S \geq 0 \ , \quad \text{and} \quad \bar{\beta}(1-\bar{P}) - \bar{\alpha}\bar{P} + S \geq 0 \ . \tag{3.2}$$

We wish to show that the optimal value of S is not in general zero.

To isolate the cross-subsidization market failure, we must purge the model of other factors which could result in a non-zero optimal value for S. First, since the effort decision for either accident is made when it is known that the other accident cannot occur, the seemingly unrelated events externality is inoperative, and need not concern us. Second, if the expected marginal utilities of income were different for the two projects, a cross-subsidy would be desirable to equalize the expected marginal utilities. (The cross-subsidy, in this case, is effectively a form of insurance.) To isolate the cross-subsidization market failure, we want to rule out this possibility, and so require that the expected marginal utilities of income be the same for both projects. Third, cross-subsidization can be desirable because it provides an indirect form of ex ante (before the individual has made his effort decision) randomization, when direct ex ante randomization is excluded (as in this paper).[20] The way we shall proceed is to prove that $S \neq 0$ in general at the social optimum, and then argue that this is not due to any considerations related to randomization.

From (3.1) the first-order conditions of the individual's effort decision problem are

$$(-\hat{u}_0+\hat{u}_1)\hat{P}' = 1 \quad \text{and} \quad (-\bar{u}_0+\bar{u}_1)\bar{P}' = 1 \ . \tag{3.3}$$

Substituting (3.2) into (3.3), eliminating the β's , and totally differentiating the resulting equations gives

$$\frac{\partial\hat{e}}{\partial\hat{\alpha}} = - \frac{\left(\hat{u}_0' \dfrac{\hat{P}}{1-\hat{P}} + \hat{u}_1'\right)\hat{P}'}{\dfrac{\hat{P}''}{\hat{P}'} + \dfrac{\hat{u}_0'(\alpha+S)(\hat{P}')^2}{(1-\hat{P})^2}} \qquad \frac{\partial\bar{e}}{\partial\bar{\alpha}} = - \frac{\left(\bar{u}_0' \dfrac{\bar{P}}{1-\bar{P}} + \bar{u}_1'\right)\bar{P}'}{\dfrac{\bar{P}''}{\bar{P}'} + \dfrac{\bar{u}_0'(\alpha-S)(\bar{P}')^2}{(1-\bar{P})^2}}$$

$$(3.4)$$

$$\frac{\partial\hat{e}}{\partial S} = \frac{\dfrac{\hat{u}_0'}{1-\hat{P}}\,\hat{P}'}{\dfrac{\hat{P}''}{\hat{P}'} + \dfrac{\hat{u}_0'(\alpha+S)(\hat{P}')^2}{(1-\hat{P})^2}} \qquad \frac{\partial\bar{e}}{\partial S} = - \frac{\dfrac{\bar{u}_0'}{1-\bar{P}}\,\bar{P}'}{\dfrac{\bar{P}''}{\bar{P}'} + \dfrac{\bar{u}_0(\alpha-S)(\bar{P}')^2}{(1-\bar{P})^2}}$$

where the denominators are negative in the relevant range.

Next we set up the social optimization problem by substituting (3.2) into (3.1) and treating the planner as choosing $\hat{\alpha}$, $\bar{\alpha}$, and S . The corresponding first-order conditions are

$$\hat{\alpha}: \quad (-\hat{u}_0'+\hat{u}_1')\hat{P} - \hat{u}_0' \frac{(\hat{\alpha}+S)}{1-\hat{P}}\,\hat{P}'\,\frac{\partial\hat{e}}{\partial\hat{\alpha}} = 0 \ , \tag{3.5a}$$

$$\bar{\alpha}: \quad (-\bar{u}_0'+\bar{u}_1')\bar{P} - \bar{u}_0' \frac{(\bar{\alpha}-S)}{1-\bar{P}}\,\bar{P}'\,\frac{\partial\bar{e}}{\partial\bar{\alpha}} = 0 \ , \text{ and} \tag{3.5b}$$

$$S:\quad (-\hat{u}_0'+\bar{u}_0') - \frac{\hat{u}'_0(\hat{\alpha}+S)\hat{P}'}{1-\hat{P}}\frac{\partial\hat{e}}{\partial S} - \frac{\bar{u}'_0(\bar{\alpha}-S)\bar{P}'}{1-\bar{P}}\frac{\partial\bar{e}}{\partial S} = 0\ . \tag{3.5c}$$

Note that $(-u_0'+u_1')P = \underline{u}' - u_0'$, where $\underline{u}' = (1-P)u_0' + Pu_1'$ is the average marginal utility of income. Multiply (3.5a) by $\frac{\hat{u}_0'}{\hat{\underline{u}}'}$ and (3.5b) by $-\frac{\bar{u}_0'}{\bar{\underline{u}}'}$ and add the resulting equations to (3.5c). This yields

$$\frac{dEU}{dS} = -\frac{(\hat{u}_0')^2}{\hat{\underline{u}}'} + \frac{(\bar{u}_0')^2}{\bar{\underline{u}}'} = 0\ . \tag{3.6}$$

From (3.4), it follows that giving a subsidy retards effort, while imposing a lump-sum tax stimulates it. The natural conjecture therefore is that one wants to tax policies for which, damage fixed, the probability of accident is more sensitive to the size of the subsidy, since doing so will reduce the probability of that accident and the deadweight loss associated with it substantially, while the subsidy provided on the other accident will increase its probability only slightly. This intuition is supported by (3.6). If we impose the requirement that the expected marginal utility of income for the two projects be the same, then (3.6) reduces to $\hat{u}_0' = \bar{u}_0'$. In the absence of cross-subsidization, more insurance is provided against that accident for which the elasticity of the accident probability with respect to the amount of insurance provided, <u>ceteris paribus</u>, is lower. Suppose this is accident 1. Then one expects that with $\underline{S}=0$, $\hat{u}_0' > \bar{u}_0'$ and therefore, from (3.6), that $\left.\frac{dEU}{dS}\right|_{S=0} < 0$, which implies that accident 1 should be subsidized.

The above argument can be formalized by considering a special case, in which the two accidents are the same except with respect to the probability-of-accident functions. Define $\emptyset = -\frac{\alpha P'}{1-P}\frac{\partial e}{\partial \alpha}$ (thus $\emptyset$ is the elasticity of the probability of <u>not</u> having an accident with respect to α , the amount of insurance provided), and assume that: i) $\tilde{\emptyset} < \hat{\emptyset} < 0$, where $\hat{\emptyset}$ and $\tilde{\emptyset}$ are constants when profits are zero, and ii) the probability-of-accident functions differ in such a way that $\hat{\underline{u}}' = \tilde{\underline{u}}'$ when $S=0$ in competitive equilibrium. Then (3.5a) and (3.5b), evaluated at $S=0$, become

$$\hat{\alpha}: \hat{\underline{u}}' - \hat{u}_0' + \hat{u}_0'\hat{\emptyset} = 0, \text{ and} \tag{3.5a'}$$

$$\bar{\alpha}: \tilde{\underline{u}}' - \tilde{u}_0' + \bar{u}_0'\bar{\emptyset} = 0 . \tag{3.5b'}$$

Since $\tilde{\emptyset} < \hat{\emptyset} < 0$, then $\hat{u}'_0 > \tilde{u}'_0$, which implies that $\left.\frac{dEU}{dS}\right|_{S=0} < 0$. Note that this local cross-subsidization has a non-zero <u>first-order</u> effect on expected utility.

Cross-subsidization may be desirable because it provides an indirect form of <u>ex ante</u> randomization. That is, it might be efficient to <u>ex ante</u> randomize either or both the accident 1 and accident 2 contracts, but when such direct randomization is not permitted, cross-subsidization between the accident 1 and accident 2 contracts may provide an indirect or second-best form of randomization. We wish to demonstrate that the desirability of local cross-subsidization in the above example is independent of such randomization considerations. To do this, we show that local[21] ex ante randomization of either contract has a zero first-order effect on expected utility.

Undertake any local ex ante randomization of the contract against accident 1. Arnott and Stiglitz [1988b] have shown that all the potential gains from ex ante randomization can be obtained with two contracts. Thus, without loss of generality, we may assume that the randomization entails two contracts. Contract A occurs with probability Q and makes arbitrarily small profits of $\Pi_A = \Delta$, while contract B occurs with probability (1-Q) and makes a profit of $\Pi_B = -\frac{Q\Delta}{1-Q}$. Both contracts A and B maximize expected utility subject to the respective levels of profit. Thus,

$$EU = QEU_A + (1-Q)\ EU_B\ ,$$

where

$$EU_A = u\left(w - \frac{\Pi_A + \alpha_A P_A}{1 - P_A}\right)(1-P_A) + u\ (w-d+\alpha_A)\ P_A - e_A\ ,$$

and similarly for contract B. For both contracts, individuals choose effort to maximize expected utility, which yields $e_A = e(\alpha_A, \Pi_A)$ and $e_B = e(\alpha_B, \Pi_B)$. Thus, we may write $EU_A = EU_A(\alpha_A, \Pi_A)$ and $EU_B = EU_B(\alpha_B, \Pi_B)$. Hence,

$$\frac{dEU}{d\Delta} = Q\left\{\frac{\partial EU_A}{\partial \alpha_A}\frac{d\alpha_A}{d\Delta} + \frac{\partial EU_A}{\partial \Pi_A}\frac{d\Pi_A}{d\Delta}\right\} + (1-Q)\left\{\frac{\partial EU_B}{\partial \alpha_B}\frac{d\alpha_B}{d\Delta} + \frac{\partial EU_B}{\partial \Pi_B}\frac{d\Pi_B}{dA}\right\}.$$

Now, since α_A and α_B are chosen to maximize expected utility,

$$\frac{dEU}{d\Delta} = Q\,\frac{\partial EU_A}{\partial \Pi_A} + (1-Q)\frac{\partial EU_B}{\partial \Pi_B}\left(\frac{-Q}{1-Q}\right)\ , \text{ and } \left.\frac{dEU}{d\Delta}\right|_{\Delta=0} = Q\,\frac{\partial EU}{\partial \Pi} - Q\,\frac{\partial EU}{\partial \Pi} = 0\ .$$

Since the above argument holds for any Q, it also holds for the optimal Q. Thus, whether or not local *ex ante* randomization is desirable, it has a zero first-order effect on expected utility. Since local cross-subsidization has a non-zero first-order effect on expected utility, the desirability of local cross-subsidization must be independent of randomization considerations.

Thus, *ceteris paribus*, insurance for accidents in which moral hazard is more (less) severe than "average" should be taxed (subsidized) since doing so stimulates "average" effort, thereby reducing the deadweight loss associated with moral hazard.

4. Pecuniary Externalities Market Failure

Recall from the discussion in section 1 that the market failure here arises because insurance firms fail to take into account that collectively the amount of insurance they provide affects producer prices and profits, which in turn affect individuals' effort at accident avoidance. In the classic, competitive economy, these pecuniary externalities "do not matter"--they cause transfers, but do not generate inefficiency. In economies with moral hazard, however, shadow prices deviate from market prices. With heterogeneous individuals, the marginal deadweight loss associated with an extra dollar of consumption by Mr. A may be larger than that for Mr. B. In this case, the transfers generated by pecuniary externalities can alter the aggregate deadweight loss associated with moral hazard. Thus, pecuniary externalities "matter"--affect the efficiency of the economy--when moral hazard is present.[22]

Rather than present a general characterization of the nature of the inefficiency, we shall, as in the previous two sections, develop a simple example in which the cause of market failure is transparent.

We assume that there is only one type of risk and one consumer good. Now, however, the accident repair industry does not have constant costs; instead, the cost of repairing the damage from an accident is an increasing function of the number of accidents. We assume, furthermore, that there are two classes of individuals--workers who have to drive and therefore face the risk of accident, and rentiers who sit at home consuming the profits from the accident-repair industry.[23]

The expected utility of workers is

$$EU = (1-p(e))u(w-\beta)+p(e)u(w-D'(p(e))+\alpha)-e \ , \tag{4.1}$$

where $D(p(e))$ is the total cost of damage repairs as a function of the number of accidents, with $D' > 0$ and $D'' \geq 0$. Workers pay the marginal cost, $D'(p(e))$, and choose effort, ignoring the fact that their collective effort affects the price of damage repairs. Thus,

$$p'(-u_0+u_1) = 1 \ , \text{ where } u_0 = u(w-\beta) \text{ and } u_1 = u(w-D'+\alpha) \ . \tag{4.2}$$

This implies $e = e(\alpha,\beta)$, which substituted into (4.1) gives $EU = V(\alpha,\beta)$.

Competitive insurance firms, meanwhile, choose α and β , taking (4.2) into account, but like workers ignoring the endogeniety of the cost of damage repairs. They effectively maximize expected utility subject to a zero profit constraint, i.e.

$$\max_{\alpha,\beta} V(\alpha,\beta) \quad \text{s.t.} \quad B^c(\alpha,\beta) = \beta(1-p)-\alpha p = 0 \ , \tag{4.3}$$

where B^c denotes the market budget constraint.

At this competitive equilibrium, renters obtain an income of $I = D'p - D$. Can the planner do better, given the same informational constraints as the market? Suppose the planner provides the renters with I and then chooses α and β . The consumer's choice of effort as a function of α and β is the same as before. The planner's resource constraint, however, differs from firms'. It is

$$B^o(\alpha,\beta) = \beta(1-p)-\alpha p-D(p) + D'(p)p - I = 0 \ . \tag{4.4}$$

The market's choice of α and β is characterized by

$$-\frac{V_\alpha}{V_\beta} = \frac{d\beta}{d\alpha}\bigg|_{B^c} = \frac{p+(\alpha+\beta)p'\frac{\partial e}{\partial \alpha}}{(1-p)-(\alpha+\beta)p'\frac{\partial e}{\partial \beta}} \ , \tag{4.5}$$

while the planner's is characterized by

$$-\frac{V_\alpha}{V_\beta} = \frac{d\beta}{d\alpha}\bigg|_{B^o} = \frac{p+(\alpha+\beta)p'\frac{\partial e}{\partial \alpha}-D''p'p\frac{\partial e}{\partial \alpha}}{(1-p)-(\alpha+\beta)p'\frac{\partial e}{\partial \beta}+D''p'p\frac{\partial e}{\partial \beta}} \ . \tag{4.6}$$

Since $p' < 0$, $p > 0$, $\frac{\partial e}{\partial \beta} < 0$ and $\frac{\partial e}{\partial \alpha} < 0$, while $\frac{V_\alpha}{V_\beta} > 0$,

then $\left.\frac{d\beta}{d\alpha}\right|_{B^c} > \left.\frac{d\beta}{d\alpha}\right|_{B^0}$ at the competitive equilibrium unless $D'' = 0$.

In α-β space the budget constraint perceived by the market is steeper than the real resource constraint, which implies that the market undersupplies insurance. Collectively, firms ignore that if they provide more insurance, damage costs go up, which, since this is equivalent to a fall in α , stimulates effort. Thus, the market perceives the responsiveness of effort to increased insurance to be greater than it actually is, and hence provides too little insurance.

The externality identified here could be corrected in a variety of ways. It would be internalized if both insurance and damage repairs were provided by the same company; this is an extension of the exclusivity requirement. Alternatively, the government could subsidize automobile accident insurance, which in general equilibrium with the consumer good as numeraire entails taxing repairs.

The example of this section was rather specific. The essential point is that, with moral hazard, pecuniary externalities have real efficiency effects that are ignored by the market. Thus, we term the inefficiency identified in this section the <u>pecuniary externalities market failure</u>.

5. Conclusions

For over a quarter of a century, the fundamental theorems of welfare economics, the formalization of Adam Smith's invisible hand, have been the central propositions in welfare economics. Though the informational assumptions underlying the theorems were generally not made explicit, intuitive discussions of the advantages of the market focussed on the "informational economy" of the price system. This paper examined the behavior of competitive markets under a particular informational hypothesis. We postulated that there are many misfortunes against which individuals wish to purchase insurance and the occurrence of which are affected by their actions; moreover, insurance firms recognize that these actions, though not directly observable, will be affected by the nature of the insurance coverage provided.

These moral hazard problems are pervasive in the economy. They arise not only in explicit insurance policies, but also in the implicit insurance associated with labor markets (wages not equal to the marginal revenue product), land markets (sharecropping), capital markets (with equity and loan contracts, when there is a finite probability of default which can be affected by the borrowers' actions) and product markets (product guarantees), etc. We have contended in this paper that economies in which these moral hazard problems are present contain numerous forms of potential inefficiency and are essentially never constrained Pareto efficient. Our analysis therefore casts serious doubt on the relevance of the fundamental theorems of welfare economics, and on the basic results concerning the efficient decentralizability of economies.

The presence of these externalities has both descriptive and prescriptive consequences.[24] We discuss the descriptive consequences first. As usual whenever there are externalities, there are private incentives for the internalization of those externalities. Our theory provides a rationale for the kind of interlinking of labor, land, and credit markets often observed in LDCs. (Indeed, our paper may be viewed as a generalization as the earlier Braverman-Stiglitz [1982] results in this area.) It also provides a rationale for firms to subsidize health care programs, which may reduce the losses associated with absenteeism and employer-financed health insurance. Our analysis may also provide part of the explanation for why insurance firms typically provide insurance against several different risks.

The intertemporal linkages, in particular the externalities between insurance provided at different dates, provide part of the explanation for long-term relations, for why individuals should work for the same employers for many periods, or why the same bank should provide credit over several periods. (See also Stiglitz and Weiss [1983].)

There is also the issue of the relationship between competitiveness and exclusivity. Exclusivity is not conceptually inconsistent with full competitiveness, provided each individual has perfect information regarding the full set of insurance contracts offered before he signs a contract which will cover all his risks for his entire life.[25] In fact, however, individuals are typically poorly informed at the time they sign their first insurance policy, and gradually acquire more information through costly search and by switching firms. Thus, there is a tradeoff between competitiveness and extended exclusivity. How the market will resolve this tradeoff is a difficult issue. But it seems safe to say that the market will be characterized by imperfect competition, only partial exclusivity, and constrained inefficiency.

Though we have identified a set of externalities which might lead to exclusive relations among a pair of economic agents, possibly covering a wide range of transactions (insurance covering various kinds of risks), the enforcement of these exclusive relations is frequently either costly or infeasible. As a result, in many insurance contexts, individuals obtain insurance against a particular risk from a variety of sources. For example, health insurance is provided not only by the individual's insurance firm, since, if he is sick, his employer generally gives him sick leave and his family will continue to provide him food and shelter. The aphorism "a friend in need is a friend indeed" can be translated "true friends provide insurance." There is a widespread view that a critical function of non-market institutions is to remedy the deficiencies of markets. Elsewhere (Arnott-Stiglitz, 1988d), we have shown that this need not be true; whether the supplemental insurance provided by non-market institutions is welfare-improving depends on whether these non-market insurance providers monitor the level of care provided by the insured.

The prescriptive implications of our results are somewhat more ambiguous. We would not argue that we have established an overwhelming case for government intervention wherever there is moral hazard. What we have established is that an ideal government can, through intervention, improve the performance of a market economy, as we have described it. But actual governments are not ideal. The potential market failures we have identified become actual market failures only when the benefits of government intervention exceed the costs. To establish this, it will be necessary, on the benefit side, to obtain estimates of the deadweight losses associated with the inefficiencies we have identified,[26] and on the cost side, to develop models of the public sector that capture the inefficiencies to which *it* is prone. We suspect, however, that there are some instances where government intervention may be warranted. In an earlier paper, for instance, we showed how subsidies to fire extinguishers or taxes on cigarettes may be welfare-enhancing.

In any case, the government is engaged in the provision of a variety of forms of insurance, and our analysis indicates that it should take these externalities into account in the design of public insurance programs.[27]

Our analysis can be criticized in another way as overstating the case for government intervention. Throughout the paper we ignored the possibility that individuals and firms may privately contract or organize to mitigate the moral hazard problems and to at least partially internalize the externalities we have identified. Consider, for example, the case of a construction firm whose accident insurance is experience-rated and in which workers with a hangover have a significantly greater probability of accident. Workers may collectively agree to restrict their alcohol consumption on evenings before work, realizing that failure to do so will result in higher accident insurance premiums. Even though each worker would have an incentive to renege on the agreement and even though monitoring and enforcing strict compliance would be very difficult, social disapproval directed at workers who came to the job hung over would be somewhat effective. In other contexts, where members of the insurance group were more anonymous and diffuse, there would be less compliance with agreed-to safety standards.

The possibility that individuals may privately cooperate to mitigate moral hazard and to partially internalize the externalities we have identified raises the question: In this context, what advantages does government have over coalitions of individuals? One is _universality_, along with which come advantages of scale and scope in reaching, monitoring, and enforcing agreements. For example, it is considerably cheaper to have the police monitor reckless driving than it would be for each insurance company to monitor its own clients' driving. Relatedly, it is much cheaper for the government to decide on universal safety standards than it is for each firm to reach agreement with its own workers on safety standards. A second advantage the government has is _the power to tax_. Suppose that as a result of moral hazard, people smoke too much. While an insurance company can "tax" smoking by making a client's premium dependent on his cigarette consumption, to do this it has to monitor the client's cigarette consumption, which is excessively costly. The government could imperfectly monitor each individual's consumption of cigarettes at lower cost, by requiring that storeowners record the identity of all cigarette purchasers. But more cost-effective than this is for it to anonymously tax the sale of cigarettes, which it can do because of universality. Because of universality and the power to tax, the government could be considerably more effective in internalizing the pecuniary externalities market failure than the collectivity of firms. A third advantage is the government's _monopoly_ _on compulsion_. Private contracting requires agreement among the parties, whereas governments can compel with no _quid pro quo_. The inter-firm transfers required to internalize the cross-subsidization market failure would not be possible without compulsion. Finally, the government can and does restrict the terms of private contracts, for instance the forms of punishments that can be meted out.

BIBLIOGRAPHY

Arnott, R., and J. Stiglitz, 1985, "Labor Turnover, Wage Structures, and Moral Hazard: The Inefficiency of Competitive Markets", Journal of Labor Economics, 3, 434-462.

Arnott, R., and J. Stiglitz, 1986, "Moral Hazard and Optimal Commodity Taxation", Journal of Public Economics, 29, 1-24.

Arnott, R., and J. Stiglitz, 1987, "The Existence and Properties of Market Equilibrium with Moral Hazard", mimeo.

Arnott, R., and J. Stiglitz, 1988a, "The Basic Analytics of Moral Hazard", Scandinavian Journal of Economics, forthcoming.

Arnott, R., and J. Stiglitz, 1988b, "Randomization with Asymmetric Information", Rand Journal of Economics, forthcoming.

Arnott, R., and J. Stiglitz, 1988c, "Price Equilibrium, Efficiency, and Decentralizability in Insurance Markets," mimeo, Stanford University.

Arnott, R., and J. Stiglitz, 1988d, "Dysfunctional Non-market Institutions and the Market," NBER Working Paper 2666.

Arrow, K., Aspects of the Theory of Risk Bearing (Helsinki: Yrjo Jahnssönin Säätio, 1965).

Braverman, A. and J. Stiglitz, 1982, "Sharecropping and the Interlinking of Agrarian Markets," American Economic Review, 72, 695-715.

Greenwald, B., and J. Stiglitz, 1986, "Externalities in Economies with Imperfect Information and Incomplete Markets", Quarterly Journal of Economics, 101, 229-264.

Hellwig, M., 1983a, "Moral Hazard and Monopolistically Competitive Insurance Markets", University of Bonn, discussion paper no. 105.

Hellwig, M. 1983b, "On Moral Hazard and Non-Price Equilibria in Competitive Insurance Markets", University of Bonn, discussion paper no. 109.

Marshall, J.M., 1976, "Moral Hazard", American Economic Review, 66, 880-890.

Newbery, D. and J. Stiglitz, 1987, "Wage Rigidity, Implicit Contracts, Unemployment and Economic Efficiency," Economic Journal, 97, 416-430.

Pauly, M., 1974, "Overprovision and Public Provision of Insurance," Quarterly Journal of Economics, 88, 44-61.

Shavell, S., 1979, "On Moral Hazard and Insurance," *Quarterly Journal of Economics*, 93, 541-562.

Spence, M., 1978, "Product Differentiation and Performance in Insurance Markets", *Journal of Public Economics*, 10, 427-447.

Stiglitz, J., 1988, "Technological Change, Sunk Costs, and Competition," *Brookings Papers on Economic Activity*, 3, 1987, Special Issue on Macroeconomics, M.N. Baily and C. Winston (eds.), The Brookings Institution, pp. 171-177.

Stiglitz, J. and A. Weiss, 1981, "Credit Rationing in Markets with Imperfect Information," *American Economic Review*, 71(3), 393-410.

Stiglitz, J. and A. Weiss, 1983, "Incentive Effects of Terminations: Applications to the Credit and Labor Markets," *American Economic Review*, 73(5), 912-927.

FOOTNOTES

* This is the third of a series of papers investigating competitive equilibrium when insurance markets are characterized by moral hazard. Arnott and Stiglitz [1988a] examines the behavior of both the insurer and the insured, showing that, as a result of moral hazard, neither indifference curves nor feasibility sets, in general, have the usual convexity properties. Arnott and Stiglitz [1987] analyzes the existence and properties of equilibrium in insurance markets with moral hazard.

Earlier versions of this paper were presented at the L.S.E., the Norwegian School of Economics and Business Administration, M.I.T., C.O.R.E., and Dortmund, Yale, and Columbia Universities. We are indebted to participants in those seminars and to Bruce Greenwald for helpful discussions.

Financial support from the National Science Foundation, the Hoover Institution, the Olin Foundation, and the Social Science and Humanities Research Council of Canada is gratefully acknowledged.

1. Moral hazard-incentive problems also arise in imperfect capital markets, even when individuals are not risk-averse. See Stiglitz-Weiss [1981]. The welfare economics for that case is similar to that presented here.

2. These events, moreover, occur in many states of nature. This, by itself, presents no serious problem.

3. Complete insurance equalizes marginal utilities of income across states of nature. So long as accidents do not affect the marginal utility of income (at any income level), providing complete insurance eliminates all incentives for accident avoidance. For a more complete analysis, see Arnott and Stiglitz [1988a] and [1987].

4. Because these externalities operate exclusively through the price system, they are sometimes referred to as pecuniary.

5. The welfare economics of the other case are treated in Arnott and Stiglitz [1988c]. We use the terms "observable" and "monitorable" interchangeably, and whenever we use either we assume verifiability. More generally, it should be clear that these are distinct concepts; an action may be observable by the two parties to a contract, but not verifiable by a third party, and therefore not legally enforceable. Enforcement in such circumstances may rely on reputation mechanisms. For a discussion of this distinction and its implications, see Newbery and Stiglitz [1987].

6. This is equivalent to the individual purchasing insurance from more than one company, subject to the requirements that each insurance company write its insurance conditional on the insurance that the individual purchases from all other companies.

7. Under these circumstances, the equilibrium insurance contract maximizes the utility of the insured subject to the non-negativity constraint on profits, and is therefore evidently efficient.

8. This paper ignores adverse selection effects in order to identify the particular market inefficiencies which arise from moral hazard. It should be clear that many of the inefficiencies which we identify here also relate to markets with adverse selection.

9. Different levels of damage associated with the same type of accident are treated as separate events.

10. Alternatively, we could have assumed that the government pays out these profits and taxes as a lump-sum payment to individuals. See footnotes 12 and 13 for how the analysis is modified.

11. It is perhaps unrealistic to assume that there are forms of insurance for which the government can specify price but not quantity, but doing so simplifies the analysis. The government may be able to indirectly control the price through tax-subsidy instruments, but incorporating these indirect controls complicates the analysis without changing the qualitative results.

12. If tax revenues plus profits, R, are distributed to individuals as lump-sum payments, then (1.16) remains unchanged, but the derivative of V now contains a term, the derivative of V with respect to income, times the derivative of R.

If the government can impose lump-sum taxes, the feasibility constraint (1.13) is dropped. The market equilibirium is still not Pareto efficient. (See footnote 13.)

13. In the case where profits and tax revenues are rebated to individuals, and lump-sum taxes can be imposed, the equation corresponding to (1.16) takes on exactly the same form, but now λ has the interpretation of the expected marginal utility of income.

The derivatives will take on different values, because of the induced income effects. The derivatives are now general equilibrium derivatives; that is, for each value of, say, η (and all the other parameters of the model) we calculate the general equilibrium solution; as η changes, each of the variables characterizing the equilibrium changes; the magnitude of the change is given by the general equilibrium derivative.

14. The two accidents could be fire and death, and the activity "care in smoking in bed".

15. A decrease in automobile insurance may induce individuals to drive more attentively, making them sufficiently more tired when they arrive at their destination that they are more likely to fall asleep while smoking in bed (that is, the marginal disutility of the effort required to undertake the fire-accident avoidance activity is increased.)

16. Walking round the house at night checking that windows are locked (to prevent burglary) facilitates checking that the wood-stove door is closed, elements turned off, etc. (to prevent fire).

17. The first-order condition for e_2 is

$$[(1-P_1)(u_0-u_2) + P_1(u_1-u_3)](-P_2') - 1 = 0.$$

The increase in insurance against accident 1 stimulates e_2 if it causes the term in square brackets to increase. Since β_1 increases, u_0-u_2 increases, but since α_1 increases too, u_1-u_3 decreases, and which effect dominates depends on whether the marginal utility of income falls more or less rapidly as income increases.

18. The reason why we stress this is that if these state variables are observable, then the terms of a policy will be made contingent on the current values of the stock variable, and there is no externality. For an important example where this distinction is clearly significant, see Arnott and Stiglitz [1985].

19. In the present model, where all individuals are identical, induced changes in the demands for various goods may result in changes in producer prices, and hence in profits. But any loss in welfare as a consumer from an increase in a price is exactly offset by a gain in welfare as a shareowner.

20. In an earlier paper, we showed that equilibrium insurance contracts may, under not restrictive conditions, be characterized by both <u>ex ante</u> and <u>ex post</u> randomization. (Arnott and Stiglitz [1988b]).

21. If this condition is satisfied, a <u>large</u> amount of <u>ex ante</u> randomization may still be desirable. See Arnott and Stiglitz [1988b].

22. It is generally true that pecuniary externalities matter in economies with distortions. The unobservability of effort, which gives rise to moral hazard, may be viewed as a distortion when individuals are risk-averse.

23. We may imagine that the consumer goods industry produces a single consumer good with constant returns to labor. The price of the consumer good is the numeraire, and the wage is determined as the marginal product in this industry. In the car repair industry, labor is combined with some other factor provided by rentiers. Each worker's labor supply is inelastic.

24. This paper has focused on those cases where the quantity of insurance purchased is observable. When it is not, similar inefficiencies arise; in addition, however, there is a further inefficiency associated with the quantity of insurance purchased, which may be partially remedied by imposing a tax on insurance purchases. This point is elaborated in Arnott and Stiglitz [1988c].

25. The argument that markets may be efficient even with a single insurance firm has much of the flavor of the contestability arguments that all that is required for competitive-like outcomes (efficiency, zero profits) is <u>potential</u> competition. As Stiglitz [1988] has shown, even the slightest sunk costs alter this conclusion. Markets will not be efficient, and profits will not be zero.

26. Since adverse selection and moral hazard almost invariably appear together, before such estimation is possible, it will be necessary to develop the welfare economics of moral hazard <u>cum</u> adverse selection.

27. Adverse selection plays an important role, both in understanding the institutional structure of markets involving risk and in designing appropriate policies. It should be noted that adverse selection gives rise to a set of market failures analogous to those we identified in the paper as stemming from moral hazard. For example taxation of commodities and cross-subsidization between insurance policies can be employed to partially relax self-selection constraints. Furthermore, as here, pecuniary externalities alter the efficiency loss associated with adverse selection. These and other points are developed in Greenwald and Stiglitz [1986].

Abstract

This paper shows that, except in certain limiting cases, competitive equilibrium with moral hazard is constrained inefficient. The first section compares the competitive equilibrium and the constrained social cptimum in a fairly general model, and identifies six types of market failure. Each of the subsequent sections focuses on a particular market failure.

Arnott/Stiglitz: Welfare Economics
February, 1989