

NBER WORKING PAPER SERIES

INFORMATION AND ECONOMIC
EFFICIENCY

Richard Arnott
Bruce Greenwald
Joseph E. Stiglitz

Working Paper No. 4533

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November, 1993

Paper prepared for presentation at the New Orleans Meetings of the American Economic Association, January 1992. Financial support of the National Science Foundation and the Hoover Institution is gratefully acknowledged. This paper is part of NBER's research program in Public Economics. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

INFORMATION AND ECONOMIC
EFFICIENCY

ABSTRACT

Is an economy with adverse selection, moral hazard, or an incomplete set of risk markets "constrained" Pareto efficient? There are two sets of papers addressing this question, one asserting that, under seemingly quite general conditions, the economy is constrained Pareto efficient, the other (to which we have contributed) that it is not. In this paper, we delineate the differences in assumptions between the two sets of papers, and under our assumptions present an intuitive proof of the Pareto inefficiency of market equilibrium with moral hazard and identify what it is that the government can do that the market cannot.

Richard Arnott
Department of Economics
Boston College
Chestnut Hill, MA 02167
and NBER

Bruce Greenwald
Columbia University School of
Business
810 Uris Hall
New York, NY 10027

Joseph E. Stiglitz
Department of Economics
Stanford University
Stanford, CA 94305-6072
and NBER

Information and Economic Efficiency

Richard Arnott, Bruce Greenwald, and Joseph E. Stiglitz¹

Arrow and Debreu's proof of the fundamental theorems of welfare economics is one of the crowning achievements of economic science. Adam Smith's invisible hand conjecture — that competition would lead to efficient resource allocation — was put on a rigorous footing. The assumptions required were very strong. Their work has stimulated a stream of papers attempting to show that the results hold under weaker conditions.

One line of that research has attempted to identify the role of assumptions concerning perfect information and complete markets. Is a competitive economy with adverse selection, moral hazard, or an incomplete set of risk markets "constrained" Pareto efficient? The word constrained emphasizes that though one surely expects adverse selection and moral hazard to affect the nature of the equilibrium, the relevant question is: even when the informational imperfections cannot be corrected, can government intervention make some individuals better off without making anyone else worse off?

In recent years, this line of research has culminated in two strands of papers, one asserting that, under seemingly quite general conditions, the economy is constrained Pareto efficient, the other (to which we have

contributed) that it is not. We will not comment here on the underlying ideological bases of these alternative research programs or the potentially large differences in policy implications arising from these different views of the potential role of government.

Our objectives here are more narrow: to attempt to delineate the differences in *assumptions* going into the different results; to present a new proof of the Pareto inefficiency of market equilibrium with moral hazard under our informational assumptions, which will serve to provide a stronger intuition for why markets are constrained Pareto inefficient; and to identify what it is that the government can do that markets cannot.

I. The Basic Argument

In an earlier paper (Greenwald and Stiglitz [1986]), we showed that informational imperfections give rise to externality-like effects. For example, in models of adverse selection, where employers cannot identify the ability of an individual worker, but the average quality of workers depends on wages and prices, government might affect the quality of the labor pool by a set of selective taxes and subsidies, designed to encourage the more productive workers to supply more labor and the less productive workers less. In models of moral hazard, where an individual's level of care at accident avoidance depends both on the level of insurance coverage and prices, the government can increase the level of care by subsidizing complements to care and taxing substitutes. Thus, if smoking increases the probability of a fire, then taxing cigarettes reduces accident rates because it reduces smoking.

The fact that government actions can have effects seems obvious in both cases. Why do we argue that there are "externalities" giving rise to

potential scope for government intervention? The reason is that the individual, in making his decision, in the first case about the level of labor supply and in the second case about the level of care, focuses only on her private costs and benefits. But there are further ramifications. In the adverse selection model, workers are wage takers. When all good workers work a little more, the average productivity of those supplying labor increases. The firm takes the average quality of those it hires as given, but its profits obviously increase (at a given wage) with the average quality. Thus, the increase in labor supply of the more productive workers has an "external" effect on the profits of firms, which they collectively fail to take into account.

Or consider the moral hazard problem. The individual, in making his effort decision, knows that the insurance firm cannot monitor care. He bases his decision on the cost of care versus the cost of an accident, and the changed likelihood of an accident resulting from greater care. He will choose the level of care to maximize his expected utility. At the margin, were he to exert slightly greater care, there would be a second-order effect on his own expected utility, but because the accident probability would go down, there would be a first-order effect on the insurance company's profits; there is an externality, which he fails to take into account.

We can express this externality-like effect in another way, by "netting out" the insurance firm. Assume that premiums, determined in competitive markets, equal the actuarial value of the losses. Consider any particular level of benefits in the case of an accident. Each individual, in deciding how much care to exert, takes the premium as given. But when all individuals exert more care, the premium is reduced, and they all can be made better off. The government, by subsidizing complements to care, can

induce individuals to take greater care. If individuals have optimized their budget over various goods, there is a second-order loss from the distortion associated with the induced change in consumption patterns, but a first-order effect arising from the reduced premiums.

There are four basic sets of cases in which government intervention might not be able to help. The first is where there is only one good. Then the government obviously cannot use commodity taxes to attain a Pareto improvement. It was natural for many of the early papers exploring the role of adverse selection, moral hazard, or incomplete markets to focus on single commodity models; after all, the particular phenomenon under study had nothing *directly* to do with many commodities. Many of the papers in the strand of literature attempting to generalize Arrow and Debreu fell into the "trap" of not realizing that the welfare properties of one-commodity models are special. (See, for instance, Diamond [1967] and Shavell [1979].)

The second set of cases involve those where the seeming market imperfection is just that, only a "seeming" imperfection. The unobservability of effort at accident avoidance would not cause a problem if the elasticity of effort with respect to insurance benefits were zero (i.e. individuals would always take the efficient level of care, regardless of whether they were being watched). This seems — and is — a trivial result. But equally trivial (but perhaps slightly less transparent) is the result that if everyone is identical, the absence of risk markets has no consequences, since even if there were risk markets, there would be no trade on them. Thus, one cannot use representative agent models to study the consequences of the absence of risk markets. We will not list the papers that have fallen into this trap.

A third set of cases involves imposing a set of constraints on the

government which is artificial, but under which the market can do as well as the government. These cases have arisen primarily in the context of the analysis of market equilibrium with incomplete risk markets, where it was not apparent what the appropriate comparison ought to be. Diamond suggested that in evaluating the stock market, the constraint to put on the government was that it make payments to individuals as linear functions of profits (since share ownerships make distributions linear functions of profits). As we have seen, in general with more than one commodity, even with this restrictive constraint, equilibrium is not constrained Pareto efficient. But there is one case where it is; if everyone has unitary price elasticities of demand, then (if output is the source of randomness) profits (price times quantity) will be constant. Hence, the stock market provides no risk transfer at all, and constraining government to make payments a linear function of profits is constraining them too to making no risk transfer.

Grossman [1977] set out explicitly to find the constraint under which the market allocation was efficient. Following Samuelson, he reasoned, in effect, that surely there must be some constrained optimization problem for which the market represented the solution. The resulting notion he called Social Nash Optimality (SNOP); it entailed (in a two-period model) first fixing the level of transfers to each individual in each state of nature. The problem was that this was a completely artificial constraint: to be sure, in any particular equilibrium, there were particular levels of transfers. But the question was, if the government intervened, say by taxing investment in some sector, why should it take the transfers as fixed? Diamond's constraint was a far more natural one, though still somewhat artificial since with options and futures markets more complicated (yet still simple) payment structures can be made.

In the first set of cases, the government could not affect the commodity composition, since there was only one commodity. In the final set of cases, the government cannot affect the commodity composition any better than the market because the market can directly control patterns of consumption. That is, if, say, in the case of moral hazard, the insurance firm could directly observe consumption bundles, it could make insurance coverage conditional on individuals consuming particular consumption bundles. While the insurance firm could not monitor the level of care, it could clearly do as well as any government which could use only tax policy to control consumption patterns, and thereby, indirectly, affect the level of care.

Actually, to attain constrained Pareto optimality, the private insurance firms have to do more than that; they must constrain the amount of insurance that individuals purchase from outsiders, either directly (by not allowing such purchases), or indirectly, by making the premiums charged depend on the amounts purchased from outsiders. Outside purchases of insurance have externality-like effects, since with the additional purchases, individuals have less incentive to exercise care, and therefore the profits of *all* those providing insurance are affected deleteriously.

In short, with effectively complete control of the quantities of all commodities consumed and all insurance purchased, the government could do no better than the private market.

This case is obviously not of much economic interest; insurance firms can, at most, monitor an individual's purchases of goods not his consumption. If there were more than one type of individual in the economy, then the "forced" levels of consumption associated with the optimal insurance contract would entail differing marginal rates of

substitution across individuals, and there would, accordingly, be an incentive to trade. Indeed, the possibility of such arbitrage is one of the reasons that limited use of non-linear pricing (taxes) is made, except for utilities, like electricity and telephone service, where resale is difficult.

While insurance firms do make some attempt to monitor the purchases of insurance coverage for the same accident, these are imperfect; for example, health insurance firms do not monitor the implicit insurance provided by employers and family (Arnott and Stiglitz [1991]). And they do not monitor the insurance provided against other risks, which may affect the level of care significantly (e.g. because of cross elasticities generated by a common pool of savings). Health insurance firms do not even attempt to monitor the most relevant aspects of consumption — such as alcohol and tobacco — and the notion that they might extend this to the entire range of commodities purchased by individuals seems ludicrous.

Yet the assumption of such monitoring lies behind the Prescott-Townsend [1984] arguments for the efficiency of market economies with moral hazard. Discussions of their paper have focused on their use of lotteries, and such lotteries do play an important role in the analysis of existence; the non-convexities that are inherently associated with moral hazard (Arnott and Stiglitz [1988b]) and other information problems give rise to existence problems (Arnott-Stiglitz [1990a]) in the absence of randomization, and provide incentives within market contexts to randomize (Arnott and Stiglitz [1988a]). The issue we are emphasizing here is not the desirability of such randomization, nor its practical relevance. Rather, the issue is what are the "objects" over which randomization occurs? Can the contracts completely specify consumption quantities? Or are the contracts more restrictive, e.g. specifying payments under certain contingencies, or

possibly even levels of consumption of some commodities, but not of *all* commodities?

In short, if the private market has *enough* information, it can do without *complete* information; complete knowledge of an individual's consumption may suffice, at least for obtaining a constrained Pareto optimum, when direct knowledge about her actions is not obtainable. But information is more incomplete than Prescott and Townsend envisage. Government can affect care *indirectly* through affecting consumption patterns, in ways which private firms cannot. Before turning to why this is so, we briefly sketch a new proof showing how taxation can lead to a Pareto improvement. Our proof highlights the role of self-selection (or incentive compatibility) constraints.

II. A New Proof of the Constrained Pareto Inefficiency of the Market with Moral Hazard

We present a simplified case with moral hazard, where: i) all individuals are identical; ii) each has the possibility of two levels of care, e^H and e^L , with p^i the accident probability when the care level is e^i , with $p^H < p^L$ (H represents high care, L low care); iii) either a fixed-damage accident occurs, or it does not; iv) an individual's purchases of insurance are observable; v) production possibilities are linear, and goods are measured such that the producer prices of all goods equal unity; and vi) the utility function is separable between consumption and effort.

The market solution in the absence of commodity taxation is that contract $\{\alpha, \beta\}$, where α is the net benefit in the event of accident and β the premium, which maximizes expected utility subject to the zero profit constraint

$$\alpha p(e(\alpha, \beta)) = \beta(1 - p(e(\alpha, \beta))),$$

where effort is a function of the benefit and premium. An individual's expected utility, with contract $\{\alpha, \beta\}$, when he chooses care level e^i , is $V(\alpha, \beta; e^i)$. Individuals choose high care with little insurance, and low care with a lot. The switch line is defined as those $\{\alpha, \beta\}$ for which the individual is indifferent between the two care levels:

$$V(\alpha, \beta; e^H) = V(\alpha, \beta; e^L).$$

The switch line as well as the zero profit constraint are depicted in Figure 1. The actual level of utility attained with contract $\{\alpha, \beta\}$ is defined by $\max_e V(\alpha, \beta, e) = V^*(\alpha, \beta)$. The indifference curves, given by $V^*(\alpha, \beta) = \bar{V}$, are also drawn in Figure 1. The market equilibrium which maximizes V^* subject to the zero profit constraint, is at Θ . We focus on the case where Θ is at the intersection of the switch line and the zero profit locus for the high-care activity (defined by $\alpha p^H = \beta(1 - p^H)$). This provides the most extensive insurance coverage possible that is consonant with non-negative profits and with the individual undertaking the safe activity.

A "small" amount of *commodity taxation* has two effects. First, it has a "second-order" effect from the standard deadweight loss which it induces, i.e. the loss of consumer surplus increases with the square of the tax rate. Secondly, it may shift the switch line out, as illustrated in Figure 2. Because at Θ , the indifference curve is not, in general, tangent to the zero profit line for the safe activity, any outward movement in the switch line has a first-order effect on utility. *Any tax or subsidy which makes the safe activity (high care) slightly more attractive (at θ) than the risky activity allows the provision of additional insurance. Since insurance was strictly rationed, the gain from the provision of the additional insurance outweighs the second-order effect of the loss of consumer surplus.*

To see this formally, we generalize our model to allow for commodity taxation. The vector of goods is X , and the associated vector of consumer prices, q . Let V^H be expected utility when activity H is undertaken. The indirect utility function is now $V^H(\alpha, \beta, q)$. With commodity taxation, the social revenue constraint is²

$$\alpha p^H = \beta(1-p^H) + (q-1)X^H$$

where

$$X^H = p^H X^{H1} + (1-p^H)X^{H0},$$

where X^{H1} is the consumption vector in the event of an accident and X^{H0} in the event of no accident. The problem of the government is to choose $\{\alpha, \beta, q\}$ to maximize V^H subject to the social revenue constraint and the self-selection constraint that the individual expend high care, i.e. $V^H \geq V^L$.

We write the Lagrangian

$$\begin{aligned} \mathcal{L} = & V^H + \gamma(\alpha p^H - \beta(1-p^H) - (q-1)X^H) \\ & + \lambda(V^H(\alpha, \beta, q) - V^L(\alpha, \beta, q)). \end{aligned}$$

To see that corrective taxation is welfare-enhancing, we differentiate the Lagrangian, evaluate the derivative at $q = 1$ — that is, at a situation where there is no differential taxation — and show that in fact the derivative is not, in general, equal to zero, so that some differential commodity taxation is desirable.

²In the absence of commodity taxation, the market solution and the constrained social optimum coincide. Here the government both provides insurance and taxes commodities. Thus, strictly speaking, what we show is that a Pareto improvement is possible if the government both taxes commodities and takes over the provision of insurance from the market. When the government taxes commodities and the market provides insurance, the analysis is more complicated but taxation remains desirable.

Differentiation of the Lagrangian with respect to q, α , and β yields³

$$\frac{\partial \mathcal{L}}{\partial q_k} \Big|_{q=1} = V_k^H - \gamma X_k^H + \lambda (V_k^H - V_k^L)$$

$$\frac{\partial \mathcal{L}}{\partial \alpha} \Big|_{q=1} = V_\alpha^H + \gamma p^H + \lambda (V_\alpha^H - V_\alpha^L) = 0$$

$$\frac{\partial \mathcal{L}}{\partial \beta} \Big|_{q=1} = V_\beta^H - \gamma(1-p^H) + \lambda (V_\beta^H - V_\beta^L) = 0$$

where V_k^i denotes $\frac{\partial V^i}{\partial q_k}$. Now $V_k^i = -p^i \mu^{i1} X_k^{i1} - (1-p^i) \mu^{i0} X_k^{i0}$ (by

Roy's Identity), $V_\alpha^i = p^i \mu^{i1}$, and $V_\beta^i = -(1-p^i) \mu^{i0}$, where μ^{i0} is the marginal utility of income with care level i in the event of no accident, etc.

Then multiplying $\frac{\partial \mathcal{L}}{\partial \alpha} \Big|_{q=1}$ by X_k^{H1} , $\frac{\partial \mathcal{L}}{\partial \beta} \Big|_{q=1}$ by $-X_k^{H0}$, and adding the

resulting equations to $\frac{\partial \mathcal{L}}{\partial q_k} \Big|_{q=1}$ yields

$$\frac{\partial \mathcal{L}}{\partial q_k} \Big|_{q=1} = \lambda \left[p^L \mu^{L1} (X_k^{L1} - X_k^{H1}) + (1-p^L) \mu^{L0} (X_k^{L0} - X_k^{H0}) \right].$$

Defining $\bar{\mu}^L = p^L \mu^{L1} + (1-p^L) \mu^{L0}$, and

$\bar{X}_k^i = \frac{p^L \mu^{L1} X_k^{i1} + (1-p^L) \mu^{L0} X_k^{i0}}{\bar{\mu}^L}$ to be "average" consumption of good

k with care level i , this equation can be rewritten as

$$\frac{\partial \mathcal{L}}{\partial q_k} \Big|_{q=1} = \lambda \bar{\mu}^L (\bar{X}_k^L - \bar{X}_k^H)$$

³The terms $\partial X/\partial q$ etc. disappear since they are preceded by $(q-1)$ which equals zero with $q = 1$.

$$< 0 \quad \text{if} \quad \bar{X}_k^H > \bar{X}_k^L$$

$$> 0 \quad \text{if} \quad \bar{X}_k^H < \bar{X}_k^L$$

Thus, unless "average" consumption of the good is the same in the two activities, it always pays to tax or subsidize the commodity. Again, the intuition behind this result is straightforward. The amount the individual is hurt or helped by taxing a commodity is proportional to the amount of the good he consumes. Thus, if the amount he consumes of some commodity should he undertake the safe activity is less than the amount he consumes of that commodity should he undertake the risky activity, then a tax on that commodity will hurt him less if he is undertaking the safe activity. At the margin, this encourages him to use the safe activity; if he was indifferent before, now he strictly prefers the safe activity. Conversely, if the amount he consumes of some commodity with the safe activity exceeds that with the risky, then a subsidy on that commodity benefits him more when he undertakes the safe activity; if he was indifferent before, now he again prefers the safe activity.⁴

⁴This result is not due to nonconvexity. With a continuum of activities and convex indifference curves, an analogous result obtains. And since randomization may be desirable only when indifference curves are nonconvex, the result is not due either to our having restricted the analysis to deterministic contracts and taxes. If there are welfare gains from a small randomization, they are second-order (see Arnott and Stiglitz [1988a]). Since the welfare gains here for a small tax are first-order, they cannot be related to randomization.

III. The Powers of Government

In recent years, there has developed a folk wisdom that anything the government can do, the private sector can do — and probably better. This ignores the fact that government has powers — powers of compulsion and proscription — which no private firm has.⁵ Private economic relations are by their nature voluntary (ignoring the Mafia). The government has the power to impose a tax on *all* production of a commodity. It also has the power to proscribe unauthorized production of a commodity, a power which makes the enforcement of its taxing powers easier. Contrast the problem of a firm wishing to discourage smoking or to encourage the use of smoke detectors with that of the government. The firm could subsidize smoke detectors, but it would find it difficult to prevent the individual from reselling the subsidized smoke detector to someone else. The firm might "insist" on its customer buying all cigarettes at above-market prices through the firm store (effectively imposing a tax), but how could it enforce such a requirement? General equilibrium economists might suggest that the firm subsidize all other commodities, but the practical problems of doing this should be obvious; besides, this does not resolve the problem of resale. The government can impose a tax or subsidy on *all* production, something that no single firm can.

The externality-like effects which we have identified are far more pervasive than the problems that were discussed in the older "market failures" literature. For instance, externalities that arise in the provision of insurance in the labor market (for instance, against the risk that a worker finds out after he has joined a firm that the idiosyncratic match is bad) reach out from each

⁵For a further articulation of these issues, see Stiglitz [1990].

firm to virtually all other firms in the economy. (See Arnott and Stiglitz [1985])

Our analysis has, in fact, understated the potential scope for government intervention, for we have focused on identifying instances where government can achieve a Pareto improvement only through the imposition of linear commodity taxes. There are instances where more active government intervention, e.g. through the control of quantities (say the quantity of insurance made available) can increase welfare even further. (See Arnott and Stiglitz [1990b]).

Conclusions

With incomplete markets and imperfect information, markets are not, in general, constrained Pareto efficient. From this perspective, the genius of Arrow and Debreu was to find the *almost* singular example of an economy — involving complete markets and perfect information — for which the market is efficient. The subsequent literature has identified a number of other exceptional cases in which the economy is constrained Pareto efficient. But the generic result, on the inefficiency of markets, appears robust.

Whether intervention by an actual government, rather than the ideal (though information-constrained) government we have considered, would be welfare improving, and what forms of government intervention are most likely to be welfare improving,⁶ are issues beyond the scope of this short paper.

⁶Greenwald and Stiglitz [1986] relate welfare-enhancing tax interventions to empirically observable parameters.

References

- Arnott, R. and Stiglitz, J.E., "Labor Turnover, Wage Structure and Moral Hazard: The Inefficiency of Competitive Markets," Journal of Labor Economics, Vol. 3, No. 4, October 1985, pp. 434-462.
- _____, "Moral Hazard and Optimal Commodity Taxation," Journal of Public Economics, Vol. 29, 1986, pp. 1-24.
- _____, "Randomization with Asymmetric Information," Rand Journal of Economics, Vol. 19, No. 3, Autumn 1988a, pp. 344-362.
- _____, "The Basic Analytics of Moral Hazard," Scandinavian Journal of Economics, Vol. 90, No. 3, 1988b, pp. 383-413.
- _____, "The Welfare Economics of Moral Hazard," in Risk, Information and Insurance: Essays in the Memory of Karl H. Borch, Henri Loubergé, ed., Kluwer Academic Publishers (Norwell, MA), 1989, pp. 91-122.
- _____, "Equilibrium in Competitive Insurance Markets with Moral Hazard," Working Paper, 1990a.
- _____, "Price Equilibrium, Efficiency, and Decentralizability in Insurance Markets," Working Paper, 1990b.
- _____, "Moral Hazard and Non-Market Institutions: Dysfunctional Crowding Out or Peer Monitoring?" American Economic Review, Vol. 81, No. 1, March 1991, pp. 179-190.
- Debreu, G., Theory of Value, New York, John Wiley, 1959.
- Diamond, P.A., "The Role of a Stock Market in a General Equilibrium Model with Technological Uncertainty," American Economic Review, Vol. 57, 1967, pp. 759-776.
- Geanakoplos, J., Magill, M., Quinzii, M. and Drèze, J., "Generic Inefficiency of Stock Market Equilibrium when Markets are Incomplete," Journal of Mathematical Economics, Vol. 19, 1990, pp. 113-151.

- Greenwald, B. and Stiglitz, J.E., "Externalities in Economies with Imperfect Information and Incomplete Markets," Quarterly Journal of Economics, May 1986, pp. 229-264.
- _____, "Pareto Inefficiency of Market Economies: Search and Efficiency Wage Models," AEA Papers and Proceedings, Vol. 78, No. 2, May 1988, pp. 351-355.
- Grossman, S.J., "A Characterization of the Optimality of Equilibrium in Incomplete Markets," Journal of Economic Theory, Vol. 15, 1977, pp. 1-15.
- Newbery, D. and Stiglitz, J.E., "The Choice of Techniques and the Optimality of Market Equilibrium with Rational Expectations," Journal of Political Economy, Vol. 90, 1982, pp. 223-246.
- _____, "Pareto Inferior Trade," Review of Economic Studies, Vol. 51, 1984, pp. 1-12.
- _____, "Wage Rigidity, Implicit Contracts, Unemployment and Economic Efficiency," Economic Journal, Vol. 97, 1987, pp. 416-430.
- Pauly, M., "Overprovision and Public Provision of Insurance," Quarterly Journal of Economics, Vol. 88, 1974, pp. 44-61.
- Prescott, E.C. and Townsend, R.M., "Pareto Optima and Competitive Equilibria with Adverse Selection and Moral Hazard," Econometrica, Vol. 52, 1984, pp. 21-45.
- Shavell, S., "On Moral Hazard and Insurance," Quarterly Journal of Economics, Vol. 93, 1979, pp. 541-562.
- Stiglitz, J.E., "On the Economic Role of the State," in The Economic Role of the State, A. Heertje, ed. (Bank Insinger de Beaufort NV), Basil Blackwell, 1990.