NBER WORKING PAPER SERIES

WHEN CAN PARTIAL PUBLIC INSURANCE PRODUCE PARETO IMPROVEMENTS?

Amy Finkelstein

Working Paper 9035
http://www.nber.org/papers/w9035

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2002

I am grateful to Daron Acemoglu, Peter Diamond, Bengt Holmström, Ben Olken, and Jim Poterba for helpful comments. The views expressed herein are those of the author and not necessarily those of the National Bureau of Economic Research.

When Can Partial Public Insurance Produce Pareto Improvements?
Amy Finkelstein
NBER Working Paper No. 9035
July 2002
JEL No. H42, D82

ABSTRACT

Wilson (1977) provided the striking result that the government can always Pareto dominate a pooling equilibrium in a private insurance market with adverse selection by providing the pooling policy as a compulsory public policy and allowing individuals to buy supplementary private insurance. I show that this Pareto improving role for the government does not derive from its unique capacity to compel participation in a public insurance program. Rather, it stems from the fact that, with the introduction of the public policy, individuals may now hold multiple insurance policies: one public and one private. If, instead, we relax the assumption of the Wilson model that individuals may only hold one private insurance policy, the private market equilibrium is always second best Pareto efficient and there is no possibility of Pareto improvement through government intervention. Whether in fact individuals are restricted to purchasing only one private insurance policy – and hence whether there is scope for Pareto improvement through government policy in this model – varies in a predictable manner across different insurance markets.

Amy Finkelstein
National Bureau of Economic Research
1050 Massachusetts Ave.
Cambridge, MA 02138
afinkels@nber.org

1. Introduction

A ubiquitous form of government intervention in insurance markets is compulsory public provision of partial insurance coverage, with the option for individuals to buy supplemental insurance on the private market. Social Security systems in countries around the world are one such example; many provide partial annuitization for the elderly, who can supplement this with additional private annuities. Public disability insurance programs in the U.S. and in Canada, national health insurance programs in Canada and in the U.K., and the U.S. Medicare system are further examples of partial public insurance programs that permit private supplementation.

One economic rationale for such government intervention in insurance markets is the inefficiency of private insurance markets when there is adverse selection. Wilson (1977) first noted that when the private market equilibrium with adverse selection is a pooling equilibrium, government intervention that provides a compulsory partial insurance policy at the population-average actuarially fair price and allows individuals to purchase supplemental coverage on the private market can always be Pareto improving. This analysis has been widely cited in support of the welfare-enhancing potential of various partial public insurance programs such as mandatory Social Security programs with supplementary private annuity markets (see for example Eckstein et al. 1985 or Eichenbaum and Peled 1987) and compulsory partial health insurance coverage with supplementary private health insurance markets (see for example Diamond 1992, Feldman et al. 1998, or Neudeck and Podczeck 1996).

This paper explores the mechanism behind the government's ability to create a Pareto improvement over the private market equilibrium in the Wilson model of adverse selection in insurance markets. I demonstrate that the Pareto improvement does not stem from the government's unique capacity to compel participation in a public insurance program. Rather, the gain from introducing public insurance in the Wilson model comes from the fact that the model restricts individuals to holding only one private insurance policy. The introduction of the public program loosens this constraint by introducing the potential for individuals to hold multiple insurance policies (one public and one private). If, instead, we relax the single private policy assumption and allow individuals to hold multiple *private* insurance policies, the private market equilibrium will always be second best Pareto efficient and the introduction of a partial public insurance policy can only redistribute along the Pareto frontier from the low risk to the high risk.

Section two reviews the intuition behind the Wilson result. Section three establishes the two results of the paper. First, I show that the power of compulsion is not needed for the government to be able to Pareto improve on the Wilson equilibrium. Second, I establish that when multiple private policies are allowed in the Wilson model, the private market equilibrium is always second-best Pareto efficient. Section four considers the empirical validity of the critical Wilson assumption that consumers may purchase only a single private insurance policy. I show that the validity of this assumption varies in a predictable manner across different insurance markets. An implication is that, in the Wilson model, there is potential for Pareto improvements from public provision of health insurance, such as that provided through the Medicare program, but not from publicly-provided annuities such as those provided through defined benefit Social Security programs. Section five concludes.

2. Partial Public Insurance Provision and Pareto Improvements in the Wilson model

The standard analysis of adverse selection in insurance markets involves competitive insurance markets and two types of individuals, high risk (H) and low risk (L). These individuals differ only in their (privately known) accident probabilities, which are denoted by p_H and p_L respectively, with $0 < p_L < p_H < 1$. It is well-known that in this setting, the private market outcome will not be Pareto efficient. A useful concept therefore in markets with adverse selection is the set of second-best Pareto efficient allocations. These allocations maximize a weighted sum of the utilities of the two types, subject to an aggregate resource constraint and incentive compatibility constraints for both types (Crocker and Snow 1985). A variety of equilibrium concepts have been proposed for competitive insurance markets with adverse selection. Crocker and Snow (1985) show that some, such as those of Rothschild and Stiglitz (1976), Grossman (1979), Riley (1979), and Wilson (1977), do not necessarily produce even second-best Pareto efficient outcomes, while others, such as that Miyazaki-Spence (Miyazaki (1977), Spence (1978)), are always second-best Pareto efficient.

When the private market equilibrium is not second-best Pareto efficient, it is interesting to consider whether and how the government can Pareto improve on the private market equilibrium. Wilson (1977) established that when the private market equilibrium is a pooling equilibrium, the government can *always* achieve a Pareto improvement.[1]

Wilson introduced a new "foresight" equilibrium concept defined as follows:

Definition: A Wilson equilibrium is a set of policies in which:

i) consumers choose a single insurance policy to maximize expected utility
ii) each policy earns non-negative profits individually
iii) there is no other set of policies outside of the equilibrium set which, if offered, would earn positive profits in the aggregate and non-negative profits individually, *after the unprofitable policies in the original set have been withdrawn.*[2]

He showed that the resultant equilibrium will exist, will be unique (except in a knife-edge case), and will be either a separating or a pooling equilibrium. These two types of equilibria are illustrated in Figure 1. The vertical and horizontal axes represent income in the states with and without an accident respectively. The point E represents the individual's endowment with no insurance. Points on the 45 degree line represent points of full insurance. Movements to the northeast represent increasing utility. The line HE (LE) represents the set of policies that earn zero expected profits when high (low) risk individuals buy them. Line EF (the market odds line) represents the set of policies priced at the population-average actuarially fair price. In a pooling equilibrium, the amount of insurance is determined by the tangency of the low risk type's indifference curve (U_L) with the market odds line EF; it is denoted by γ in the figure; it provides less than full insurance. In a separating equilibrium, which is identical to that of Rothschild and Stiglitz (1976), the high risk type pays his own-type actuarially fair marginal price and gets full insurance (at a_H) while the low risk type pays his own-type actuarially fair marginal price and gets exactly that amount of insurance that leaves the high risk type indifferent between purchasing his equilibrium policy and deviating toward purchasing the low risk policy (at a_L); this is less than full insurance. The equilibrium is a separating one if the low risk type prefers a_L to γ and a pooling one if, as drawn, the low

[1] When the equilibrium is separating, there may or may not be room for Pareto improvement. This was first noted by Rothschild and Stiglitz (1976) and discussed in more detail by Crocker and Snow (1985). I concentrate on the pooling case here because it is where the result is more dramatic – the government can *always* enact a Pareto improvement – and where the Wilson model offers something different from other equilibrium concepts.

[2] The italicized portion in the above definition indicates the modification of the equilibrium concept used by Rothschild and Stiglitz (1976).

risk type prefers γ to a_L.[3]

Wilson shows that if the government provides and requires everyone to purchase the private market pooling equilibrium policy γ and allows supplemental policies to be purchased, the result is a Pareto improvement over the initial pooling equilibrium γ, This is shown in Figure 2. The government policy translates the endowment point from E to γ. We can now look for a private market equilibrium for supplemental policies from this new endowment point. γ. Line $\gamma H'$ ($\gamma L'$) represents the set of policies that earn zero expected profits when high (low) risk individuals buy them starting from the new endowment γ; line γF represents the set of policies that earn zero expected profits when both types purchase them. It is straightforward to see that the private market equilibrium cannot be pooling. By definition of the original pooling equilibrium γ, low risk types do not wish to purchase any additional insurance at the market-odds actuarially fair price. This ensures the existence of a separating equilibrium with the usual characteristics, which is shown in the figure by $\{\beta_H, \beta_L\}$. The high risk type receives the amount of supplemental insurance necessary to be fully insured and pays his actuarially fair price for the supplemental insurance; the low risk type receives the maximal amount of supplemental insurance at his actuarially fair price that can maintain incentive compatibility for the high risk type. Each of these supplemental policies earns zero profits as required in equilibrium. Moreover, since at the original pooling equilibrium individuals receive less than full insurance – and individuals therefore on the margin want to pay their actuarially fair price for more insurance – both types are strictly better off at the final allocation $\{\beta_H, \beta_L\}$ than they were at the original pooling equilibrium γ. This can be seen in the figure by comparing utility of each type at the new equilibrium –given by U_L^* and U_H^* -- to utility at the original equilibrium (U_L and U_H).[4]

3. Why Can't the Private Market Achieve the Equilibrium Induced by the Public Program?

[3] The equilibrium allocation of the Wilson "foresight" equilibrium is identical to that obtained by the Grossman (1979) "dissembling" equilibrium.

[4] Eckstein et al (1985) extend the Wilson analysis to describe the entire set of compulsory government insurance policies at the market-odds price that would produce Pareto improvements over a private market pooling equilibrium. They show that any compulsory government-provided insurance at the market-odds line with supplemental private policies will be Pareto improving as long as the government provision is less than or equal to the amount of insurance in the pooling equilibrium and greater than or equal to the amount that, when the high risk individuals then purchase supplemental insurance at their actuarially fair price to achieve full insurance, their utility is equal to the amount in the original pooling equilibrium.

As discussed above, this Wilson (1977) result has been widely cited as a rationale for a variety of partial public insurance programs. But a natural question is why, when there is a private market pooling equilibrium, the private market does itself not offer the Pareto improving supplemental policies induced by the public program. The answer is that the private market does not offer the supplemental policies because the model restricts individuals to holding at most one private insurance policy.

I establish two results. First, the power of compulsion is not necessary for the government to achieve a Pareto improvement over the private market equilibrium. Second, if we modify the Wilson model to allow individuals to purchase multiple *private* insurance policies, the private market equilibrium is always second-best Pareto efficient and there is no scope for Pareto improvements through public policy.

3.1 Pareto improvements with optional partial public insurance

Proposition I: *Suppose the government offers an* optional *public policy that replicates the private market pooling equilibrium at g. It also allows individuals to buy private insurance, and continues to offer the public policy as long as it breaks even. Then the resulting equilibrium is identical to that when the publicly-provided policy is* mandatory.

Proof: To establish that the final allocation $\{\beta_H, \beta_L\}$ is still an equilibrium when the government provides the optional policy γ, consider the two possible types of private market responses. First, the private market may offer policies available only as a supplement to the government policy. In this case, we know that the supplementary policies from the new endowment point γ that result in the final allocation $\{\beta_H, \beta_L\}$ constitute an equilibrium. Second, the private market may offer an *alternative* policy to the combination of government policy γ and the private supplemental policies that result in the final allocation $\{\beta_H, \beta_L\}$. However, there is no profitable alternative policy that the private market can offer that remains profitable after all unprofitable policies are withdrawn. To see this, we must consider three possible cases. First, there is no profitable policy from the initial endowment point E that attracts both types since the low risk types prefer β_L to their most preferred point on the market-odds line (γ). Second, there is no profitable policy from the initial endowment point E that attracts only high risk types since at

β_H high risk types receive full insurance and also receive a lump sum subsidy (equal to the distance along the 45 degree line between EH and EH′).

Third, consider policies that attract only low risk types: any policy in the shaded region of Figure 2 – below the fair odds line (EL) for the low risk type, above the low risk indifference curve U_L^* through β_L, and below the high risk indifference curve U_H^* through β_L – would attract only the low risk types and earn non-negative profits. However, if the low risk types purchase a policy in this region rather than the combination of the government pooling policy γ and the private supplemental policy that gets them to β_L, the pooling policy γ now attracts only high risk types; γ thus becomes an unprofitable policy and would be withdrawn by the government. With the pooling policy withdrawn, high risk types prefer any policy in the shaded region to their supplemental policy that provides the insurance from γ to β_H; as a result, a policy in the shaded region will attract both types and hence not be profitable.

To see that high risk types prefer any policy in the shaded region to their supplemental policy alone, note first that low risk types prefer any outcome in the shaded region to a_L (shown in Figure 1), which denotes their allocation without government intervention should they choose the separating equilibrium.[5] Now consider the outcome x, which represents the policy in the shaded region with the least amount of insurance; it lies at the intersection of line EL and the indifference curve U_L^*. Since low risk types strictly prefer x to a_L and the marginal price of the two policies is the same, x must contain more insurance than a_L. The single crossing property insurance that high risk types must therefore also strictly prefer x to a_L, and since they are indifferent – by construction – between a_H and a_L, they must strictly prefer x to a_H. High risk types must therefore strictly prefer x to the supplemental policy by itself since this policy has the same marginal price as a_H but pays out less insurance. Moreover, since outcome x represents the outcome in the shaded area with the least amount of insurance, and since low risk types at least weakly prefer all outcomes in the shaded area to outcome x, the single crossing property ensures that high risk

[5] This follows by transitivity. By construction, low risk types (weakly) prefer any outcome in the shaded region to β_L. We have already established that they prefer β_L to γ, and that they prefer γ to a_L, since the original equilibrium without government intervention is a pooling one. Thus by transitivity they prefer any outcome in the shaded region to a_L.

types strictly prefer any outcome in the shaded area to outcome x, and thus any outcome in the shared area to the supplemental policy by itself. Therefore there exists no profitable deviation to attract low risk types that will not also attract high risk types once the unprofitable pooling policy is withdrawn. ¦

The preceding argument establishes that even *optional* government provision of the market equilibrium pooling policy γ produces a Pareto improvement over the market equilibrium γ. The intuition for why compulsion is not needed stems from the fact that the pooling policy γ is priced at the population-average actuarially fair price and therefore neither risk type equates marginal rates of substitution with marginal rates of transformation. Moreover, the pooling policy provides less than full insurance. Therefore, both types prefer to buy additional insurance at their own-type actuarially fair price. By offering the optional pooling policy γ and allowing individuals to buy *supplementary* private insurance policies, the government can make both types better off than they are with just the pooling policy γ since the supplementary private market equilibrium will be a separating one which allows both types to buy some additional insurance at their actuarially fair price.

Thus the ability of the government – unlike a private firm – to compel purchase of an insurance policy is not critical to the government's ability to achieve a Pareto improvement over a private market pooling equilibrium. Indeed, the preceding argument has established that were a private firm to offer the policy γ -- and, like the government, only withdraw it if it became unprofitable – then, if individuals are allowed to purchase multiple private policies – the resulting equilibrium would be the same as that obtained through the compulsory public insurance policy γ. However, in the Wilson model, the private market cannot offer such supplemental policies to the private policy γ because it is assumed (see Wilson 1977, p.170) that individuals may only purchase a single private insurance policy. Therefore what government provision accomplishes in the Wilson model is the possibility of multiple insurance policies (one government, and one private) and thus the Pareto improvement over the pooling equilibrium.

The critical difference between the government and the private market in this model is that private firms will behave *strategically* in determining the optimal level of the pooling policy whereas the level of the government policy was established by assumption. Therefore, the private market equilibrium when individuals are allowed to buy multiple policies will not necessarily involve the final allocation $\{\beta_H, \beta_L\}$ because the equilibrium amount of cross-subsidization of high risk types by low risk types may be different than the amount provided by the policy γ. I now consider in more detail the nature of the private market equilibrium when individuals are allowed to buy multiple policies.

<u>3.2 The Wilson equilibrium with multiple private policies</u>

Proposition II: *A Wilson equilibrium modified to allow individuals to purchase multiple private policies will always be second best Pareto efficient.*

Proof: Crocker and Snow (1985) prove that the Miyazaki-Spence equilibrium is always second-best Pareto efficient. We prove that the Wilson equilibrium with multiple private policies produces the identical allocation to the Miyazaki-Spence equilibrium.

The Miyazaki-Spence equilibrium modifies the original single-policy Wilson equilibrium defined above by relaxing assumption (ii) to allow firms to offer policies that do not individually break even, as long as the set of policies offered by the firm breaks even in the aggregate. Let l denote the proportion of the population that is type H. An insurance contract can be described by the premium (P_i) that type i pays in either state of the world and the amount of the indemnity payment (q_i) that type i receives in the event of the accident. We denote by v_i the expected utility of type i. Spence (1978) showed that the Miyazaki-Spence equilibrium is the solution to the following Program I:

$$\max\ v_L(P_L, q_L)$$
s.t.
$$(1)\ lP_H + (1-l)P_L = lp_H q_H + (1-l)p_L q_L$$
$$(2)\ v_H(P_H, q_H) \geq v_H(P_L, q_L)$$
$$(3)\ v_H(P_H, q_H) \geq \max_q v_H(p_H q, q)$$

The Miyazaki-Spence equilibrium maximizes the expected utility of low risk types subject to the constraints that in aggregate the firm breaks even across all its policies (constraint 1), that the high risk type (weakly) prefer his chosen policy to that chosen by the low risk type (constraint 2), and that the high risk type is at least as well off as the best he can do purchasing insurance at his actuarially fair marginal price of p_H (constraint 3).

The original Wilson equilibrium replaces the aggregate break-even constraint (1) above with a per-policy break-even constraint for the separating or pooling contract as follows:

$$\text{(1a) } P_i = p_i q_i \text{ for } i \in \{H, L\} \quad \text{(separating equilibrium)}$$
or
$$\text{(1b) } P = \left(l p_H + (1-l) p_L\right) * q \quad \text{(pooling equilibrium)}$$

Now consider the modification to allow individuals to purchase multiple private policies.[6] Since each policy must individually break even, the set of contracts purchased must consist of combinations of separating and pooling policies. Without loss of generality, we can simplify the analysis of multiple policies by considering at most one separating and one pooling policy, since there are no gains to be had by a firm from subdividing the separating or pooling component of the contract. We can therefore write the Wilson equilibrium with multiple policies as the solution to the following Program II:

$$\max \ v_L(P_L, q_L)$$
s.t.
$$(1') \ P_i = m(p_i q_i) + (1-m)\left(l p_H + (1-l) p_L\right) * q \ \text{ for } i \in \{H, L\} \text{ and } m \in [0,1]$$
$$(2') v_H(P_H, q_H) \geq v_H(P_L, q_L)$$
$$(3') \ v_H(P_H, q_H) \geq \max_q v_H(p_H q, q)$$

The resource constraint $(1')$ requires that any allocation be produced by a convex combination of a separating and a pooling policy. Since the objective function and constraints $(2')$ and $(3')$ in Program II are identical to their counterparts in Program I, it remains to show that the resource constraints (1) and $(1')$ are equivalent. Comparing them, we can see that constraint $(1')$ imposes the additional constraint – beyond that imposed by the aggregate resource constraint (1) – that any cross-subsidization is from the low risk type to the high risk type. However, this constraint is not binding as it is already imposed by

[6] It is important to note that insurance companies' ability to charge an increasing marginal price, which is a necessary condition for maintaining a separating equilibrium, does not require that individuals purchase only one private policy. An increasing marginal price requires only that insurance companies are able to monitor total insurance purchased by any customer (Rothschild and Stiglitz 1976).

constraint (3). Thus the allocation in a Wilson equilibrium with multiple policies is identical to that in a Miyazaki-Spence equilibrium and is therefore second-best Pareto efficient. ¦

The intuition for why allowing multiple policies guarantees that the Wilson equilibrium is second best Pareto efficient is clear. The separating equilibrium of Rothschild and Stiglitz (1976) or Wilson (1977) may not be second best Pareto efficient because it does not allow for cross-subsidization of high risk types by low risk types. Such cross-subsidization weakens the incentive compatibility constraint on the high risk type and thus allows the low risk type to purchase additional insurance. This clearly makes the high risk type better off and may make the low risk type better off if his marginal utility from additional insurance exceeds his marginal disutility of the transfer to the high risk type. The Miyazaki-Spence equilibrium is always second best Pareto efficient because it permits such cross-subsidization. By allowing for a pooling policy, the single-policy Wilson equilibrium can also achieve cross-subsidization. However, a pooling equilibrium can never by itself be second best Pareto efficient because neither risk type equates marginal rates of substitution with marginal rates of transformation, and both types have less than full insurance.

Once multiple policies are allowed in the Wilson model, the model can achieve both the optimal amount of cross-subsidization through the pooling policy as well as the marginal efficiency requirements through the supplementary separating policies. The Wilson equilibrium with multiple policies will therefore always involve a separating contract; it may also involve a pooling contract in addition to the separating contract if some amount of cross-subsidization of the high risk type by the low risk type is optimal for the low risk type. Unlike the government, which can choose an arbitrary amount of cross-subsidization, the Wilson equilibrium with multiple policies results in the unique amount of cross-subsidization that maximizes the utility of the low risk type.

Since the Wilson equilibrium is second best Pareto efficient when the purchase of multiple private policies is allowed, and since the government is typically assumed to have no more information than private firms, publicly provided insurance – indeed any government intervention – becomes purely a redistributive rather than efficiency-enhancing endeavor. Clearly the government cannot redistribute further to the low risk types along the second-best Pareto frontier since the private market equilibrium achieves the low risk type's maximum utility along this frontier. But by providing a compulsory policy at the market-odds price that provides *more* insurance than the pooling policy in the multiple-policy private market equilibrium, and allowing the private market to supplement *this* public policy, the government can move along the second-best Pareto frontier to redistribute further from the low risk types to the high risk types; indeed, it can achieve any second-best Pareto efficient allocation in which the high risk type is fully insured and the low risk type has less than or equal to full insurance.

4. In which insurance markets – and why – is the single policy assumption reasonable?

The ability of a partial public insurance program to create a Pareto improvement over a pooling equilibrium therefore depends critically on whether individuals can purchase multiple private insurance policies. In practice, whether the assumption of single insurance policies – and thus the welfare-enhancing role of government – is a reasonable one varies across different types of insurance markets.

Individuals are not covered by multiple private health insurance policies. For example, in the 1996 National Health Interview Survey – which collects information on up to four different health insurance plans – only about 5 percent of adults with private health insurance report that they are covered by more than one private health insurance plan.[7] In contrast, individuals tend to hold multiple life insurance policies and multiple private annuity policies. For example, Cawley and Philipson (1999) report over one-quarter of elderly individuals with term life insurance policies hold multiple such policies. LIMRA (2000) reports that, in 1997, over half of individuals who purchased immediate, fixed payment, individual annuities in that year reported owning other private individual annuities.

Four features of the private insurance market are likely to influence whether individuals tend to hold multiple private insurance policies. First, policies that make relatively frequent payments – such as health insurance or annuities – are likely to benefit from economies of scale in administering the insurance and therefore are more likely to have individuals only hold one such policy. Life insurance policies, which will make at most one payment, do not have the same savings on administrative costs from the provision of only one policy. Second, policies that provide a service in addition to a financial payment are likely to benefit from non-financial economies of scale; this mitigates against the appeal of holding multiple

[7] Based on author's calculation.

policies. For example, HMO policies which join the provision of health services with the financial reimbursement are likely to provide services such as centralization of individual's medical records and establishment of doctor-patient relations that would make holding multiple HMO policies unattractive to the consumer. Third, the insured may be more eager to hold multiple policies when they are written as long-term contracts – such as life insurance and annuities – rather than as annual contracts such as health insurance; by holding multiple policies from multiple sources, the insured can diversify the risk in a long-term contract of any given company becoming insolvent and hence unable to pay a claim many years hence.

Fourth –and perhaps most importantly – a restriction to single policies may be more likely when moral hazard effects of insurance are large. In such circumstances, insurance companies may include co-payments and deductibles to try to strike the optimal trade-off between reducing risk exposure and creating perverse incentives; they would not want this to be undone by the purchase of additional insurance to cover these cost-sharing provisions. Moral hazard effects are likely to be largest for health, and smallest for annuities. Cutler and Zeckhauser (2000) review the empirical literature on the elasticity of demand for medical care with respect to its price and conclude that moral hazard is a quantitatively important effect of health insurance. While there is no empirical evidence of moral hazard effects of annuities, it seems likely that moral hazard is less of an issue in this insurance market than in others. Philipson and Becker (1998) present a theoretical argument for how the receipt of an annuity may lead some individuals to devote additional resources to life-extension. In practice, however, in developed countries like the United Kingdom, Finkelstein and Poterba (2000) suggest that the fraction of their income that annuitants receive from annuities is too small to make moral hazard effects likely to be quantitatively important in private annuity markets.

It is therefore not surprising that health insurance policies – which feature all four factors likely to limit buyers to one policy – are almost exclusively owned singly whereas life insurance and annuity policies – which feature at most one factor each – are commonly held in multiples. The theoretical results suggest that this difference between health insurance and annuity markets has critical implications for the potential welfare-enhancing role of government-provided partial insurance coverage. In particular, the results suggest that partial public health insurance programs such as the U.S. Medicare system and the Canadian and British national health insurance systems may potentially produce Pareto improvements along the lines suggested by the Wilson (1977) model. In contrast, defined benefit Social Security systems, which provide partial annuitization around the world, cannot produce Pareto improvements through alleviation of adverse selection pressures in the Wilson (1977) model. Whether in practice partial public health insurance programs do achieve this theoretical possibility is an open question and one that I examine empirically elsewhere (Finkelstein 2002). For a restriction to single policies is necessary but not sufficient for their to be potential for Pareto improvement through partial public insurance; we have seen, for example, that there is no scope for Pareto improvements through government intervention in the Miyazaki-Spence single policy equilibrium.

5. Conclusion

This paper has explored the mechanism behind the much-cited Wilson (1977) finding that publicly-provided partial insurance at the population-average fair odds price combined with private supplementary policies can always Pareto dominate a private market pooling equilibrium in an insurance market with adverse selection. The government's capacity for affecting this Pareto improvement stems not from its unique capacity to compel participation in an insurance program, but merely from its introduction of the potential for individuals to hold multiple insurance policies. Once individuals are allowed to hold multiple *private* insurance policies, the private market equilibrium will always be second-best Pareto efficient and government intervention therefore cannot be efficiency-enhancing. A comparison across insurance markets indicates that, in predictable ways, the single-policy assumption – and hence the potential for Pareto improvement through government intervention in this model – applies in some insurance markets but not in others.

An important avenue for further research is to explore the precise *reason* why individuals purchase only a single insurance policy, for this may effect the welfare consequences of introducing partial public insurance programs. For example, if large moral hazard effects are the primary reason for the restriction to single private health insurance policies, the social welfare gains from introducing a partial public health insurance program may be mitigated. As noted by Pauly (1974), the introduction of a partial public insurance program in a model with moral hazard instead of adverse selection can be welfare *reducing* because of the moral hazard subsidy of the private residual insurance by the public program. Incorporating such moral hazard effects into a welfare analysis of the introduction of a partial public program in the Wilson adverse selection model becomes critical if moral hazard is a primary reason for the restriction to single policies in this model

References

Cawley, John and Tomas Philipson. 1999. An Empirical Examination of Information Barriers to Trade in Insurance. American Economic Review 89(4), pp.827-846.

Crocker, Keith and Arthur Snow. 1985. The Efficiency of Competitive Equilibria In Insurance Markets with Asymmetric Information. Journal of Public Economics, 26, pp.207-219.

Cutler, David and Richard Zeckhauser. 2000. The Anatomy of Health Insurance. In Handbook of Health Insurance, Culyer AJ and Newhouse JP (eds.)

Diamond, Peter. 1992. Organizing the Health Insurance Market. Econometrica, (60): 1233-1254.

Eckstein, Zvi, Martin Echeinbaum and Dan Peled. 1985. Uncertain Lifetimes and the Welfare Enhancing Properties of Annuity Markets and Social Security. Journal of Public Economics, 26, pp.303-326.

Eichenbaum, Martin and Dan Peled. 1987. Capital Accumulation and Annuities in an Adverse Selection Economy. The Journal of Political Economy, (95): 334-354.

Feldman, Roger, Carlos Escribano and Laura Pellise. 1998. The role of government in health insurance markets with adverse selection. Health Economics (7): 659-670.

Finkelstein, Amy and James Poterba. 2000. Adverse Selection in Insurance Markets: Policyholder Evidence from the U.K. Annuity Market. NBER Working Paper No. 8045.

Finkelstein, Amy. 2002. The Interaction of Partial Public Insurance Programs and Residual Private Insurance Markets: Evidence from the U.S. Medicare Program. NBER Working Paper 9031.

Grossman, Hirschel. 1979. Adverse selection, dissembling, and competitive equilibrium. Bell Journal of Economics, pp. 336-343.

LIMRA. 2000. Immediate Annuity Buyer Study. LIMRA International, Windsor CT.

Miyazaki, Hajime. 1977. The rate race and internal labor markets. Bell Journal of Economics, pp.394-418.

Neudeck, Werner and Konrad Podczeck. 1996. Adverse selection and regulation in health insurance markets. Journal of Health Economics, (15): 387-408.

Pauly, Mark. 1974. Overinsurance and Public Provision of Insurance: The Roles of Moral Hazard and Adverse Selection. Quarterly Journal of Economics (88): 44-62.

Philipson, Tomas and Gary Becker. 1998. Old-Age Longevity and Mortality-Contingent Claims. Journal of Political Economy. (106): 551-573.

Riley, John. 1979. Informational equilibrium, Econometrica 47, 331-359.

Rothschild, Michael and Joseph Stiglitz. 1976. An Essay on the Economics of Imperfect Information. Quarterly Journal of Economics 90 (November), pp.629-649.

Spence, Michael. 1978. Product differentiation and performance in insurance markets. Journal of Public Economics, (10): 427-447.

Wilson, Charles. 1977. A model of Insurance Markets with Incomplete Information. Journal of Economic Theory 16, pp.167-207.

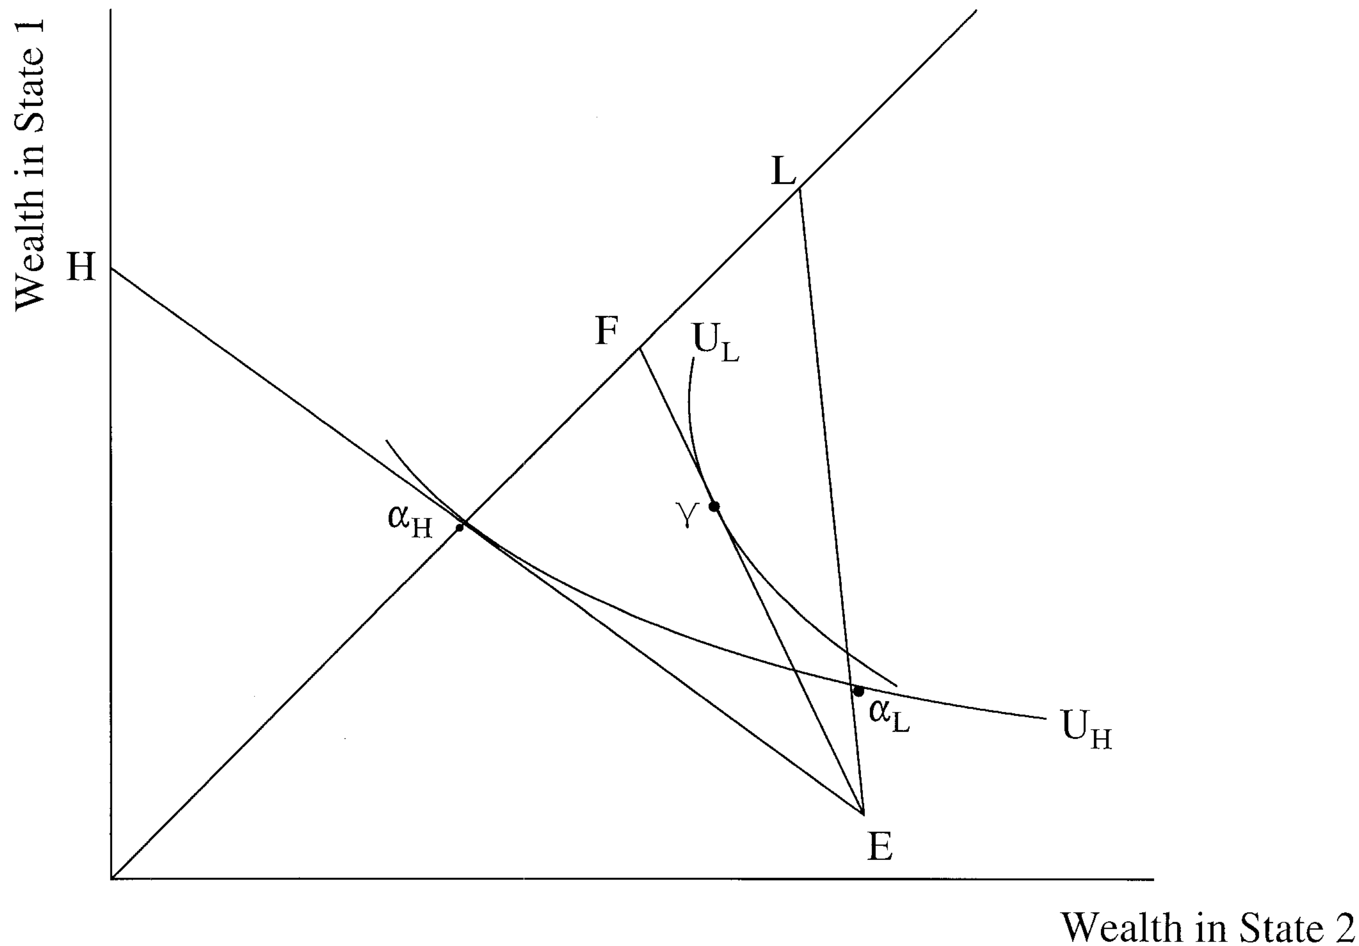

Figure 1: The Wilson equilibrium

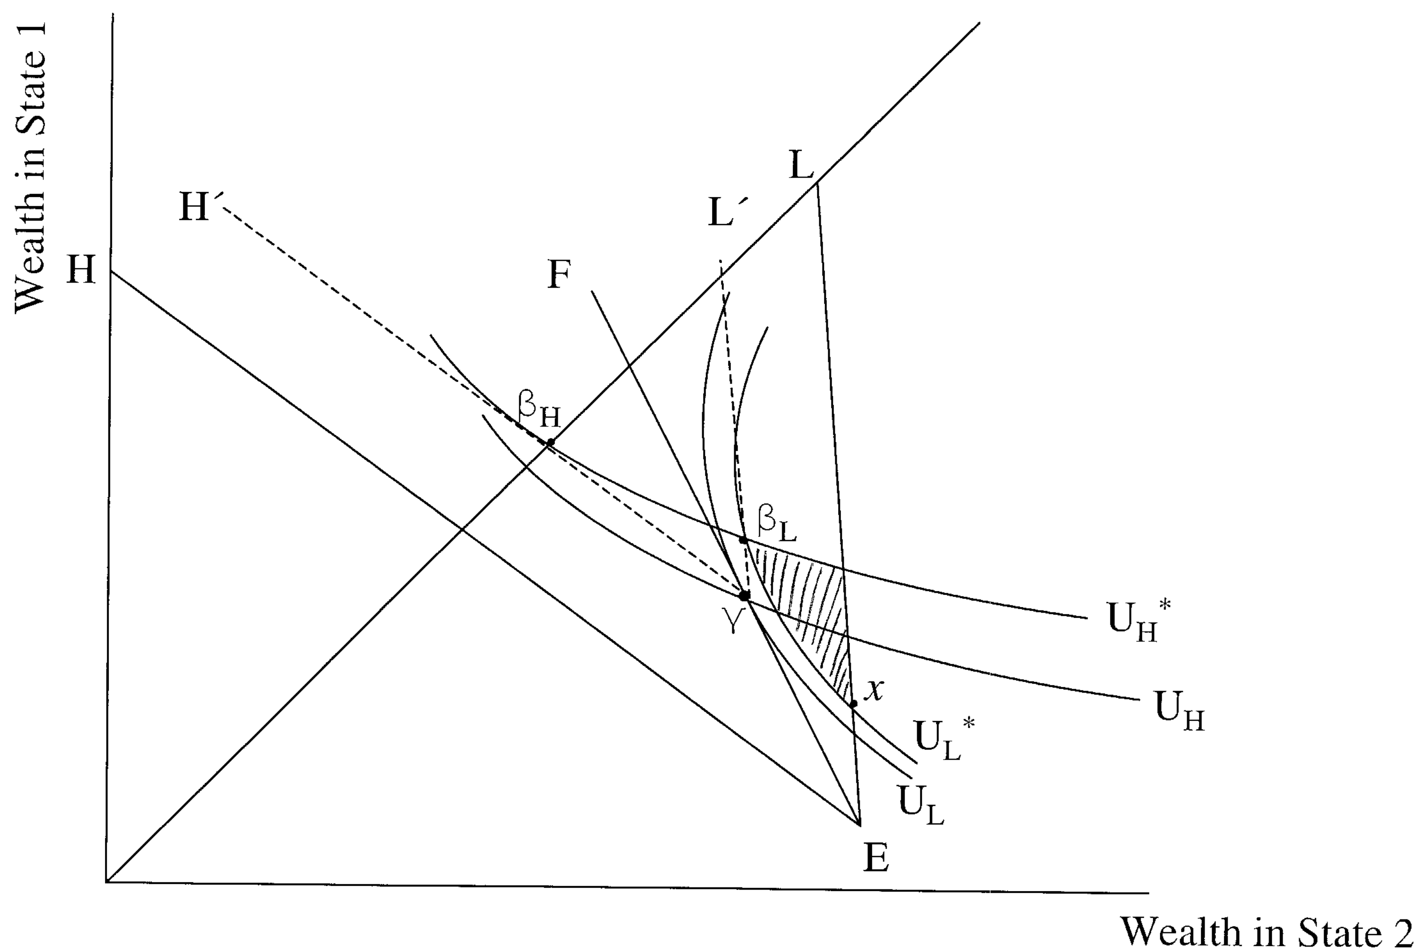

Figure 2: Pareto improvements on the Wilson equilibrium