



University of New England

**Graduate School of Agricultural and Resource Economics
&
School of Economic Studies**

**A Positivist Approach to Pigouvian Taxes based on an
Evolutionary Algorithm**

by

Phil Simmons and Oscar Cacho

No. 2000-8 – May 2000

Working Paper Series in

Agricultural and Resource Economics

ISSN 1442 1909

<http://www.une.edu.au/febl/EconStud/wps.htm>

Copyright © 2000 by Phil Simmons and Oscar Cacho. All rights reserved. Readers may make verbatim copies of this document for non-commercial purposes by any means, provided this copyright notice appears on all such copies. ISBN 1 86389 685 6

A Positivist Approach to Pigouvian Taxes based on an Evolutionary Algorithm

Phil Simmons and Oscar Cacho**

Key Words: Pigouvian tax, pollution tax, genetic algorithm, political preferences

** Phil Simmons and Oscar Cacho are Senior Lecturers in the School of Economic Studies, and members of the Graduate School of Agricultural and Resource Economics at the University of New England. Contact information: School of Economic Studies, University of New England, Armidale, NSW 2351, Australia. Email: psimmons@metz.une.edu.au.

A Positivist Approach to Pigouvian Taxes based on an Evolutionary Algorithm

1. Introduction

The Pigouvian tax is one of a class of taxes designed to mitigate against adverse welfare affects of pollution from industrial processes. It is distinguished from other pollution taxes inasmuch as it is levied on industrial output rather than on contaminants or effluent directly. In conventional treatments, such as Randall (1987) and Hartwick and Olewiler (1998) the Pigouvian tax, which may be fixed or ad-valorem, is set equal to the marginal social cost of damage arising from pollution. The tax allows polluting production to the point where the marginal benefits that it confers on producers and users are equal to the marginal social costs of pollution. Setting the tax at this level maximises welfare defined as the sum of producer and consumer surpluses associated with the polluting production or, in multi-period treatments, the discounted sum of these measures. Baumol (1972) has shown that providing markets are competitive, such taxes are allocatively efficient in the sense of being Pareto Optimal.

In this study, we are interested in how Pigouvian taxes are likely to be set when outcomes are uncertain. For example, if the polluting industry is farming, future prices and climate may be known only in terms of priors. Similarly, uncertainty may exist about social costs arising from pollution in the future. This type of uncertainty means setting the tax must be based on incomplete information and caution may be appropriate. In addition, our approach distinguishes between normative approaches such as those cited above, which identify a policy ideal, and positivist approaches that take into account government preferences. The study is in the spirit of the political economy literature from Peltzman (1976) and Becker (1983) however both our specification of government interest and solution technique, based on an evolutionary algorithm, are new.

A model is constructed with two industries where one industry externalises some of its costs onto the other. Private management of the industries is assumed to be risk neutral in the Subjective Expected Utility Model (SEUM) sense of constant marginal

utility on the part of managers, however government policy makers are assumed to be cautious, reflecting political concerns. The solution involves setting the tax using an evolutionary algorithm. It is based on a relatively new theory of risk behaviour from Szpiro (1997) and Cacho and Simmons (1999) where selection pressure, in our case political pressure, provides a motivation for caution in decision making.

2. A Pigouvian Model

Two industries, A and B, have the following production functions:

$$y_A = ax_A^a \quad (1)$$

$$y_B = bx_B^b - cax_A^a \quad (2)$$

where y_A and y_B are physical outputs and x_A and x_B are composite factor inputs for A and B respectively. $0 < \mathbf{a} < 1$, $0 < \mathbf{b} < 1$ and $a > 0$, $b > 0$ and $c > 0$ with c small relative to a and b . Output from A adversely affects output from B as a classical production externality.

Equity markets are assumed to be complete and without transaction costs allowing owners to diversify investment risk. Hence managers, with access to complete markets for investment risk, have constant marginal utility and are profit maximising in their business decision making. The managers' decision problems are:

$$\max_{x_A} \mathbf{p}_A = p_A e^{u_1} ax_A^a - c_A x_A \quad (3)$$

$$\max_{x_B} \mathbf{p}_B = p_B e^{u_2} (bx_B^b - cax_A^a) - c_B x_B \quad (4)$$

where \mathbf{p}_A and \mathbf{p}_B are profits, $p_A e^{u_1}$ and $p_B e^{u_2}$ are lognormally distributed output prices, c_A and c_B are factor input prices for A and B respectively and u_1 and u_2 are normally distributed independent random variables with $Cov(u_1, u_2) = 0$.

First Order Conditions (FOC) are:

$$\frac{\partial p_A}{\partial x_A} = \mathbf{a}p_A e^{u_1} x_A^{a-1} - c_A = 0 \quad (5)$$

$$\frac{\partial p_B}{\partial x_B} = \mathbf{b}p_B e^{u_2} x_B^{b-1} - c_B = 0 \quad (6)$$

and, solving FOCs for x_A and x_B :

$$x_A = \left(\frac{c_A}{\mathbf{a}p_A e^{u_1}} \right)^{\frac{1}{a-1}} \quad (7)$$

$$x_B = \left(\frac{c_B}{\mathbf{b}p_B e^{u_2}} \right)^{\frac{1}{b-1}} \quad (8)$$

Substituting (7) and (8) back to (1) and (2), equilibrium output from A and B with the externality are:

$$y_A = a \left(\frac{c_A}{\mathbf{a}p_A e^{u_1}} \right)^{\frac{a}{a-1}} \quad (9)$$

$$y_B = b \left(\frac{c_B}{\mathbf{b}p_B e^{u_2}} \right)^{\frac{b}{b-1}} - ca \left(\frac{c_A}{\mathbf{a}p_A e^{u_1}} \right)^{\frac{a}{a-1}} \quad (10)$$

The presence of the externality discourages production in industry B. One way to ‘correct’ for the externality is to find the ‘Coasian market solution’. That is, assume both industries are owned by the same group and managed to maximise joint profits. This provides incentives to mitigate against pollution from industry A, and because the externalised cost is effectively internalised by re-arrangement of ownership, joint profit from both industries is maximised. This hypothetical situation gives the maximum profit achievable from the joint resources in the industries and provides a

benchmark level of possible joint profit for policy makers to aim for in setting a pollution tax (Randall, 1987, p. 186-92).

If both industries are owned by the same group, the decision problem is:

$$\max_{x_A, x_B} \mathbf{P}_A + \mathbf{P}_B = p_A e^{u_1} a x_A^a - c_A x_A + p_B e^{u_2} (b x_B^b - c a x_A^a) - c_B x_B \quad (11)$$

with FOCs:

$$\frac{\partial(\mathbf{P}_A + \mathbf{P}_B)}{\partial x_A} = \mathbf{a}(p_A e^{u_1} - c p_B e^{u_2}) x_A^{a-1} - c_A = 0 \quad (12)$$

$$\frac{\partial(\mathbf{P}_A + \mathbf{P}_B)}{\partial x_B} = \mathbf{b} p_B e^{u_2} x_B^{b-1} - c_B = 0 \quad (13)$$

where (12) and (13) are solved for equilibrium outputs using (1) and (2):

$$y_A = a \left(\frac{c_A}{\mathbf{a}(p_A e^{u_1} - c p_B e^{u_2})} \right)^{\frac{a}{a-1}} \quad (14)$$

$$y_B = b \left(\frac{c_B}{\mathbf{b} p_B e^{u_2}} \right)^{\frac{b}{b-1}} - c a \left(\frac{c_A}{\mathbf{a}(p_A e^{u_1} - c p_B e^{u_2})} \right)^{\frac{a}{a-1}}. \quad (15)$$

Comparing (14) and (15) with (9) and (10) and noting $\mathbf{a}, \mathbf{b} < 1$, it follows y_A has decreased and y_B has increased as a result of joint ownership. The polluting industry does less polluting and the polluted industry is less polluted.

The solution to the optimal tax based on the functional forms of the production functions (1) and (2) is intractable using the usual solution technique based on FOC with the tax as a policy variable. However, an alternative approach is possible

through the market solution described above.

A Pigouvian tax is specified through modification of (3):

$$\max_{x_A} \mathbf{p}_A = p_A e^{u_1} (1-t) a x_A^a - c_A x_A \quad (16)$$

where t is an *ad valorem* tax on output from industry A. FOCs within the two industries are obtained from (16) and (4):

$$\frac{\partial \mathbf{p}_A}{\partial x_A} = a p_A e^{u_1} (1-t) x_A^{a-1} - c_A = 0 \quad (17)$$

$$\frac{\partial \mathbf{p}_B}{\partial x_B} = b p_B e^{u_2} x_B^{b-1} - c_B = 0 \quad (18)$$

From a public policy perspective, assuming a dollar has the same value whether in private or public domains, the welfare function is:

$$\text{welfare} = \mathbf{p}_A + \mathbf{p}_B + \text{trev} \quad (19)$$

where *trev* is tax revenue:

$$\text{trev} = t p_A e^{u_1} y_A = t p_A e^{u_1} a x_A^a \quad (20)$$

The FOC for the policy decision, $\frac{\partial \text{welfare}}{\partial t} = 0$, is a power function that cannot be solved easily with (17) and (18) to obtain solutions to the optimal private decision variables, x_A and x_B , and the policy variable, t . To solve for t , compare FOC from the market solution, (12) and (13), with FOC from the situation with the tax, (17) and (18). Find a value for t so that FOC in each situation are equivalent, requiring $t p_A e^{u_1} = c p_B e^{u_2}$. With identical FOC, the tax achieves the same outcome as the market solution in terms of incentives facing the two managers. It turns out that when $t = c p_B e^{u_2} / (p_A e^{u_1})$ the tax is equivalent to the market solution in terms of its impact

on managers' decision making. Thus, the optimal tax level is a simple function of relative prices prevailing between the industries.

To obtain a reduced form for the welfare function with t remaining exogenous, (17) and (18) are solved simultaneously for x_A and x_B . x_A and x_B are then substituted into (20), (16) and (4) providing reduced forms for substitution into the three elements of the welfare function, (19):

$$\begin{aligned}
 \text{welfare} = & -c_B \left(\frac{c_B}{\mathbf{b}p_B e^{u_2}} \right)^{\frac{1}{b-1}} + b \left(\frac{c_B}{\mathbf{b}p_B e^{u_2}} \right)^{\frac{a}{l-a}} p_B e^{u_2} \\
 & + c_A \left(\frac{c_A}{\mathbf{a}p_A e^{u_1} (1-t)} \right)^{\frac{1}{a-1}} + a (p_A e^{u_1} - cp_B e^{u_2}) \left(\frac{c_A}{\mathbf{a}p_A e^{u_1} (1-t)} \right)^{\frac{1}{a-1}}
 \end{aligned} \tag{21}$$

Numerical Assumptions

The evolutionary algorithm is solved numerically hence assumptions about the values of prices and coefficients are needed (Table 1). Using values from Table 1, the analytical solution for the optimal tax, $t = cp_B e^{u_2} / (p_A e^{u_1})$, indicates welfare is maximised at $t = 27$ per cent when output prices are set at mean values.

Using mean values for prices and values from Table 1, the welfare function is plotted against a range values for the tax in Figure 1. Figure 1 reveals a relatively smooth, convex surface (Figure 1).

3. Base Run for the Evolutionary Algorithm

Evolutionary algorithms have developed as a general research tool following early work by Fogel, Owens and Walsh (1966), Holland (1975) and others with good introductory texts being Goldberg (1989) and Mitchell (1997). Specifications of evolutionary algorithms vary between researchers and, arguably, there are as many evolutionary algorithms as there are interested researchers with computer coding skills. The evolutionary algorithm used here was developed over a number of years to

deal with problems in economic research. These problems tended to be large in terms of the number of coefficients to be estimated and often involved extracting signals from noisy environments. Two major constraints on development were the need to avoid premature convergence in problems with weak signals and to limit computer processing time needed to find solutions. While there are many approaches to specification of evolutionary algorithms most, but by no means all, involve the following elements: initialisation, selection, pairing, crossing and mutation. In explaining how our evolutionary algorithm works, each of these elements is discussed in the context of obtaining a baseline solution for the pollution tax problem.

The baseline solution provides a reference for later runs made under different assumptions and allows a check against the theoretical result for the optimal tax already obtained. It is undertaken assuming output prices are fixed at their means as with the analytical solution.

Initialisation

Initialisation is the seeding of the genesis population expected to evolve over the course of estimation. Forty random values for the tax, bounded below by zero to exclude the possibility of a negative tax, and above by unity, are drawn from a uniform distribution. These are substituted into the function to be maximised, equation (21), along with values from Table 1 with, as explained, prices held constant for the base run. Thus, forty random tax scenarios are generated and stored in a list.

Selection

In the selection stage, the objective function, (21), is solved for each random value of the pollution tax generated in initialisation and results are ranked from the tax scenario generating the highest welfare through to the tax scenario generating the lowest welfare. Highly ranked scenarios can be viewed as being ‘more fit’ and lower ranked ones ‘less fit’. Selection then occurs with the lowest ranked scenario being dropped from the population leaving a population of 39. Thus, selection pressure is gentle with only a few per cent of the population being selected against in each

generation.

Pairing

In the pairing stage, two scenarios are selected from the population to become pairs, or ‘parents’, of ‘offspring’ in the new generation. The first member of the pair is the one with the tax scenario ranked most highly in welfare terms. The second member of the pair is chosen using a ‘roulette wheel’ technique similar to that described in Goldberg (1987). That is, the member is chosen randomly from the remaining 38 members of the population with more highly ranked scenarios having a greater probability of being selected. Thus, the first member of the pair is the most highly ranked scenario and the second member can be any other scenario with the second most highly ranked scenario having the highest probability of selection. The list of probabilities of selection as a mate are generated from the fitness rankings as the first difference of a zero-one cumulative distribution for a list of consecutive numbers corresponding in length to the size of the population minus one.

Crossing and mutation

The pair ‘breeds’ in the sense that a new scenario is formed that has a tax level that is a randomly weighted average of the tax levels of each member of the pair. The weights applied to each member’s tax level in forming this average are generated from separate drawings from a uniform distribution with a range of $0.5 \pm x$. x can be viewed as a ‘mutation’ that occurs in the averaging of parental characteristics and means that offspring are not simple averages of parental pairs. In this study, x was set at 0.1. The new ‘offspring’ resulting from this combining of parental characteristics is added to the population restoring its size to forty. When values of t in the new generation were detected that violated $0 < t < 1$, the tax was reset at either zero or one depending whether the lower or upper bound had been violated.

This crossing technique is at variance with the genetic algorithm literature where parental characteristics are combined using a technique called ‘bit string swapping’ described in Goldberg (1987) and Mitchell (1997). ‘Bit string swapping’ is a more complicated form of crossing than averaging. It is based on an analogy to

DNA splitting in haploid life-forms such as viruses. Our own research began several years ago with traditional genetic algorithms based on ‘bit string swapping’ however, perhaps reflecting the types of problems that concerned us, we found superior results, in terms of convergence characteristics and computer processing time, with the numerical crossing approach described above.

Repeating the Loop

The new generation consists of 39 members from the preceding generation and one new, hopefully ‘fitter’, member. This generation then enters the cycle of selection, pairing, crossing and mutation to create yet another generation. The cycle is repeated, in this case for 150 generations, until the value of the tax converges and no higher level of welfare can be obtained.

The convergence path for welfare is reported in Figure 2 which shows the average level of welfare achieved across the population each generation. The convergence path for the average level of the tax is reported in Figure 3. It turned out that the average value for the tax in generation 150 was 26.9 per cent, close to the theoretically derived Pigouvian result of 27 per cent. This result was expected since suppression of noise in prices meant that the evolutionary algorithm was maximising a simple deterministic function.

4. Results with Stochastic Prices

In the base run reported above, output prices in each industry were set at their means. This was relaxed by generating random lognormal distributions for prices based on Table 1 values for u_1 and u_2 . The disturbances were incorporated into the algorithm in the following manner. Two lists of lognormally distributed random variables, corresponding to e^{u_1} and e^{u_2} were created with the same length as the number of generations in each run. The first element in each list was set at the means of the disturbances so that initialisation could occur at the mean. Similarly, the last term in each list was set to the mean so welfare achieved by the end of each run could be monitored at mean price levels. Each new generation, random terms from the list were introduced so each member in each generation faced the same two random

values for e^{u_1} and e^{u_2} . Thus, in generation one, each member's welfare function was calculated and ranked at the means, then in generation two, the two mean disturbance terms were replaced with two new random terms, then in generation three, another two new random terms were introduced and so on.

The experiment was conducted in three stages. In the first stage, e^{u_2} was suppressed by setting it to its mean value while e^{u_1} was stochastic. In stage two, the opposite occurred with e^{u_1} suppressed and e^{u_2} stochastic. In stage three both e^{u_1} and e^{u_2} were stochastic. Hence, the affects of selection on setting the tax were measured in stage one with noisy prices in the polluting industry, in stage two with noisy prices in the polluted industry and in stage three with noisy prices in both industries simultaneously.

Results from evolutionary algorithms usually converge however small differences arise in results from repetitions of experiments because of the presence of random mutations. Thus, in these experiments, the algorithm was re-run ten times for experimental stages one, two and three. In fact, the results for the convergent levels of the tax, reported for the ten repetitions in Table 2, were very similar for each stage.

Stage One Results

The running of the algorithm with a noisy stochastic term in the welfare function required many more generations to extract the signal than with a deterministic environment. Hence, all three runs were 10,000 generations in length which may have involved some harmless overkill. In the first run, with noisy prices in the polluting industry and fixed prices in the polluted industry, the value of the tax averaged 43.7 per cent which is considerably higher than the level of the tax of 27 per cent resulting from classical theory. This level of tax discriminates strongly against the polluting industry reflecting that prices in that industry are the only source of noise in the model. The path over which convergence occurred in the first run is shown in Figure 4.

Stage Two Results

In stage two, noise in the polluting industry price was suppressed and the sole source of noise was output price in the polluted industry. This resulted in an average tax level of 16.6 per cent over the ten repetitions, well below the level for the classical model of 27 per cent. Thus, the government was less protective of the polluting industry when it was making a relatively noisy contribution to welfare. The welfare stream is safer in terms of government's political survival when stable returns from industry A are emphasised over unstable returns from the polluted industry B.

Stage Three Results

Both industries contributed the same amount of noise in the welfare function in stage three with noisy output prices in both A and B. The resulting average tax level of 27 per cent was the same as for the classical treatment of the tax from Pigou. In this situation, the government has nothing to gain from favouring, or discriminating against, either industry in terms of stabilising welfare.

Discussion and Conclusions

The approach taken to management of risk in this study is from Szpiro (1997) and Cacho and Simmons (1999). It is based on quite different assumptions to von Neuman and Morgenstern's (1947) Subjective Expected Utility Model (SEUM) which is usually encountered in economic treatments of risk. However, our qualitative conclusion that setting of the tax discriminates against the industry which contributes most to noise in the welfare function is consistent with an SEUM treatment of the problem. For example, if welfare in the theoretical model was specified as $\text{Log}(\textit{welfare})$ and government assumed to be altruistic, hence ascribing Diminishing Marginal Utility (DMU) to government in its approach to welfare, similar qualitative results would have been obtained to those obtained with the evolutionary algorithm. That is, the solution would have discriminated against the relatively noisy industry and the deterministic Pigouvian solution would have been recovered when noise was symmetrically distributed between the industries. Similarly a 'market solution' where the joint owners of the two industries were assumed to have DMU rather than being

risk neutral would have discriminated against noise and also recovered the Pigouvian solution with symmetrical noise.

Yet, despite the similarities in qualitative results, the two theories, SEUM and ours, are fundamentally different. The central assumption in SEUM is that agents have DMU in either income or wealth. The assumption of DMU means that agents place relatively less value on marginal income when income is relatively high and hence sacrifice little, in utility terms, by undertaking saving or reducing debt. Conversely, when low-income states are encountered, marginal dollars are valued relatively highly in utility terms and borrowing and dis-saving become attractive. This means agents have incentives in expected utility terms for stabilising incomes and, in the framework of SEUM, results in powerful and generally robust theories about how government should deal with uncertainty.

With the evolutionary algorithm, policy engendered stability occurs for different reasons to SEUM approaches. Government wishes to set the tax to generate the highest possible level of welfare however it is faced with a probabilistic game because of the stochastic prices in the model. If prices in, say, the polluting industry are high relative to those in the polluted industry, then a low tax is appropriate to maximise welfare. However, if conditions change and the polluting industry price falls relative to the polluted industry price, then the high tax becomes welfare reducing and a political liability. Hence, government attempts to find a balance between possible outcomes that minimises the possibility of being 'caught' with the 'wrong' policy settings.

Despite the apparent observational equivalence in qualitative outcomes from the two approaches to risk, it is possible results from strategic models such as ours are intrinsically different to those from SEUM analyses in terms of their allocative efficiency. Some hint of this comes from Cacho and Simmons' (1999) use of the Separation Theorem in their strategic model of farm investment under uncertainty. They showed that 'strategic' farms, faced with selection pressure from the possibility of bankruptcy, are likely to operate below their risk-efficient frontiers and hence not Pareto Optimal. That is, when farm investments reflected fear of bankruptcy, risk-return outcomes could be dominated. There were potential strategies with

higher expected returns and the same uncertainty or reduced uncertainty with the same expected return than those strategies predicted by the model. Thus, caution was associated with a degree of inefficiency in the Pareto sense. In our study, the recovery of the original Pigouvian solution when uncertainty is symmetrical between the two industries indicates Pareto optimality can occur in these types of models. However, it is less clear whether the same result would prevail in stages one and two with asymmetric disturbances. This means that political pressure on government to perform in terms of welfare may lead to inefficient outcomes in terms of resource allocation in setting of pollution taxes. Doing the best by the community in terms of maximising welfare may not be the best political strategy. If the assumptions of our model are considered, with welfare being largely comprised of profits accruing to risk neutral owners, the government's preferred tax levels are clearly inappropriate except in the case where noise was symmetrical. Hence, a strong conclusion from this study is that a politically 'pressured' government may not provide the best policy settings for a pollution tax.

References

- Baumol, W.J. 1972, 'On taxation and the control of externalities', *American Economic Review*, 12(3), June, pp. 307-22.
- Becker, G.S. 1983, A theory of competition among pressure groups for political influence, *Quarterly Journal of Economics*, 93, pp. 371-400.
- Cacho, O.?. and Simmons, P.R. 1999, 'A genetic algorithm approach to farm investment', *Australian Journal of Agricultural and resource Economics*, 43(3), pp. 305-323.
- Goldberg, D.E. 1989, *Genetic algorithms in search, optimization and machine learning*, Addison Wesley, New York.
- Hartwick, J.M. and Olewiler, N.D. 1998, *The economics of natural resource use*, Addison-Wesley, New York, 2n edn, p. 31.
- Holland, J.H. 1975, *Adaptation in natural and artificial systems*, University of Michigan Press.
- Mitchell, M. 1997, *An introduction to genetic algorithms*, MIT Press, Cambridge.
- Peltzman, S. 1976, Toward a more general theory of regulation, *Journal of Law and Economics*, 18, pp. 745-771.
- Randall, A. 1987, *Resource economics: an economic approach to natural resources*, John Wiley & Sons, 2nd edn, p. 360.
- Szpiro, G. 1997, 'The emergence of risk aversion', *Complexity*, 2(4), pp. 31-9.
- von Neumann, J. and Morgenstern, O. 1947, *Theory of games and economic behaviour*, Princeton University Press, Princeton, Ch. 1-4.

Tables and figures

| coefficient or price assumption | value |
|------------------------------------|-----------|
| <i>a</i> | 0.5 |
| <i>b</i> | 0.5 |
| <i>a</i> | 2 |
| <i>b</i> | 2 |
| <i>c</i> | 0.15 |
| p_A | \$5 |
| p_B | \$9 |
| c_A | \$2 |
| c_B | \$3 |
| u_1 | $N(0, 1)$ |
| u_2 | $N(0, 1)$ |

Table 1: Numerical Assumptions for Coefficients and Prices

| Run # | Stage 1 | Stage 2 | Stage 3 |
|----------------|---------|---------|---------|
| 1 | 0.436 | 0.162 | 0.268 |
| 2 | 0.449 | 0.161 | 0.276 |
| 3 | 0.452 | 0.175 | 0.263 |
| 4 | 0.456 | 0.165 | 0.284 |
| 5 | 0.442 | 0.166 | 0.276 |
| 6 | 0.419 | 0.156 | 0.248 |
| 7 | 0.407 | 0.163 | 0.289 |
| 8 | 0.441 | 0.173 | 0.269 |
| 9 | 0.436 | 0.165 | 0.261 |
| 10 | 0.429 | 0.173 | 0.267 |
| Mean | 0.437 | 0.166 | 0.270 |
| Standard Error | 0.015 | 0.006 | 0.012 |

Table 2: Results for the Tax over Ten Repetitions




Figure 1: Plot of Welfare Against the Tax.




Figure 2: Convergence Path for Welfare over 150 Generations




Figure 3: Convergence Path for the Tax over 150 Generations




Figure 4: Convergence Path for Tax when Polluter's Price is Noisy