MPRA

Munich Personal RePEc Archive

Evolutionary Analysis of Climate Policy and Renewable Energy: Heterogeneous Agents, Relative Welfare and Social Network

Volker Nannen and Jeroen C. J. M. van den Bergh

Vrije Universiteit, Amsterdam

14. July 2008

Evolutionary Analysis of Climate Policy and Renewable Energy

Heterogeneous Agents, Relative Welfare and Social Network

Volker Nannen[*] and Jeroen C.J.M. van den Bergh[†]

ABSTRACT. We demonstrate how an evolutionary agent-based model can be used to evaluate climate policies that take the heterogeneity of strategies of individual agents into account. An essential feature of the model is that the fitness of an economic strategy is determined by the relative welfare of the associated agent as compared to its immediate neighbors in a social network. This enables the study of policies that affect relative positions of individuals. We formulate two innovative climate policies, namely *prizes*, altering directly relative welfare, and *advertisement*, which influences the social network of interactions. The policies are illustrated using a simple model of global warming where a resource with a negative environmental impact—fossil energy—can be replaced by an environmentally neutral yet less cost-effective alternative, namely renewable energy. It is shown that the general approach enlarges the scope of economic policy analysis.

Key words: agent-based modeling · behavioral economics · climate policy · evolutionary economics · relative welfare · social network

JEL classification: B52 · C73 · H23 · Q54

1. INTRODUCTION

The analysis of the economic impact of climate change and climate policy is dominated by neoclassical general equilibrium and growth models. Some models in this vein, which have played a prominent role in the IPCC and international policy debates, are: DICE (Nordhaus, 1991, 1994), RICE (Nordhaus and Yang, 1996), ENTICE (Popp, 2004), CETA (Peck and Teisberg, 1993), MERGE (Manne, 1992) and FUND (Tol, 1995). Kelly and Kolstad (1999) and van den Bergh (2004) present brief accounts and evaluations. Although these models have generated many insights, they do not represent the full range of possible approaches nor of the questions that can be addressed. They omit certain elements in their description of reality: out of equilibrium processes, choice between multiple equilibria

[*]*Faculty of Computer Science, Vrije Universiteit, De Boelelaan 1081a, 1081 HV Amsterdam, the Netherlands.* email: *vnannen@gmail.com*.

[†]*ICREA, Barcelona, and Department of Economics and Economic History & Institute of Environmental Science and Technology, Autonomous University of Barcelona, Edifici Cn, 08193 Bellaterra, Spain. Also affiliated with Faculty of Economics and Business Administration & Institute for Environmental Studies, Vrije Universiteit, Amsterdam, The Netherlands. Fellow of NAKE and Tinbergen Institute.* email: *jeroen.bergh@uab.es*.

(path-dependence), structural changes in the economy due to innovations, and the influence of income or welfare distribution on strategies. In addition, the available models assume representative agents, rational behavior, perfect information, and an aggregate production function. This approach allows for exact solutions, but it also limits the type of policies that can be studied. For example, they cannot study the effects of information provision, or of exemplary reward and punishment.

Here we present a model that starts from a set of alternative feasible assumptions offered by evolutionary economics (Nelson and Winter, 1982; Dopfer, 2005; Witt, 2008) and by agent-based computational economics (Tesfatsion and Judd, 2006; Levy et al., 2000; Epstein and Axtell, 1996). Evolutionary modeling has gained some popularity in economics, but most studies in this vein lack a policy dimension. This holds especially for applications of pure evolutionary game theory (Friedman, 1998). However, agent-based models applied to economics have rarely addressed public policy issues, and if they have done so, only in a way that does not fully exploit the policy potential offered by an evolutionary model (Janssen and Ostrom, 2006). The present study adds a policy angle to an agent-based evolutionary approach, focusing on the opportunities that an evolutionary system offers for policy design and analysis. This approach recognizes evolution in the economy rather than emphasizing the use of evolutionary algorithms to optimize non-evolutionary complex systems (e.g. Janssen et al., 2004).

The evolutionary agent-based model developed here addresses policy in a setting of global warming. The latter is endogenous to the model and depends on the source of energy used by agents in the model. These agents may be interpreted as national or regional authorities in charge of the energy policy of an independent economy. Global warming is assumed to have a negative effect on social welfare. The overall goal at the global level is to replace a resource with a negative impact on social welfare—fossil energy—by a neutral alternative, namely renewable energy. On an individual level the alternative comes with no economic advantage, and possibly even with a disadvantage, so that there is no incentive to adopt it. A complicating factor is that there is no central authority that can formulate and enforce a policy. Climate policies are usually based on international agreements, and compliance by countries is voluntary.

Each agent is modeled individually and agents are assigned only limited information and boundedly rational capabilities. Their objective is assumed to be to reach a high level of individual welfare. The only information available to the agents are the investment strategies and the income growth rates of their fellow agents. The agents believe that there is a strong causal link between these two variables. Since they prefer a high over a low income growth rate, they imitate the investment strategy of a fellow agent when that fellow agent realizes an income growth rate that is high relative to their own income growth rate and that of their other fellow agents. That is, they imitate an investment strategy when they believe that it causes a relatively high income growth rate. This approach is inspired by findings on relative welfare and income comparison effect of happiness or "subjective well-being" studies (e.g. Ferrer-í-Carbonell, 2005; Frank, 1987).

Imitation is never perfect. Small errors are introduced during the imitation process that lead to slightly different variants of the same investment strategy. While these errors are necessary to maintain diversity within the pool of strategies and to allow a population of imitating agents to find and converge on the individually optimal strategy, they also form a hitherto unexploited opportunity for the policy maker. If the desirable variant is given a selective advantage over the undesirable variant, the first will diffuse faster and will ultimately be used more often. For example, a policy that aims for agents to adopt

a greener investment strategy—in the sense that they invest less in fossil energy—can do so by convincing at least some agents that the greenest variants in the current pool of strategies lead to a relatively high income growth rate. As these strategies are imitated, the errors guarantee that some of the new variants will be greener still. An evolutionary policy can thus "breed" a new strategy by progressively giving a selective advantage to the most desirable variants.

As has been extensively discussed by Wilhite (2006), agent-based simulation of economic processes needs to give proper attention to the social network. Communication links between economic agents, individuals and institutions, are neither regular nor random. They are the result of a development process that is steered by geographic proximity, shared history, ethnic and religious affiliation, common economic interests, and much else. The social networks used for this study reproduce a number of stylized facts that are commonly found in real social networks: the small world effect (Erdős and Rényi, 1959), a high clustering coefficient (Watts and Strogatz, 1998), and a scale-free degree distribution (Barabási and Albert, 1999). Modeling the evolution of strategies in such complex social networks allows us to formulate economic policies that exploit the effect of social visibility on the diffusion of a desirable strategy.

The frequency with which the strategy of an agent is imitated depends on two factors: the relative welfare of an agent, as observed from the income growth rate, and the position of an agent in the social network. A policy that is aimed at increasing the selective advantage of a particular strategy can do so in two ways, namely by changing the relative welfare of an exemplary agent that uses such a strategy, or by changing the position of such an agent in the social network. We formulate policies for both alternatives, and compare their effect to that of a standard tax on fossil energy. One policy, *prizes*, increases the relative welfare of exemplary agents by awarding them a monetary prize. The other policy, *advertisement*, increases the social visibility of exemplary agents by increasing their connectivity in the social network. Policy tools that increase social visibility include, for instance, sponsorship of industrial fairs and scientific venues, and the production and distribution of educative material.

The remainder of this article is organized as follows. Section 2 presents the climate-economy model, including the evolutionary mechanisms of strategy formation and diffusion, and formulates the climate policies. Section 3 studies the convergence behavior of the resulting evolutionary process. Section 4 evaluates the climate policies using numerical simulations with the climate-economy model. Section 5 concludes.

2. The economic model

2.1. **General features of the model.** The present economic model is formulated in order to study the effectiveness of regulatory public policies when economic behavior evolves through social interactions. The approach focuses on climate policies and energy investment strategies, but the model can easily be adapted to other problems. Each agent controls an independent economy with its own supply and production. The agent formulates a strategy to invest current domestic income in different sectors. The returns for each independent economy are then calculated from standard economic growth and production functions. Some allocations give higher returns than others, and the goal of the agents is to find a strategy that can realize a high level of individual welfare.

The present model is loosely based on the influential work of W. D. Nordhaus, who published a series of general-equilibrium economic models of climate policy and global warming, starting with the DICE model (Nordhaus, 1992). From this model all economic

a, b	individual agents	k	average number of neighbors per agent
N	neighbors of an agent	$\mathcal{C}$	clustering coefficient of the network
s	investment strategy	σ	mutation variance of strategies
$\mathcal{F}$	fitness of a strategy	Q	income without global warming
Y	net domestic income	γ	income growth rate
K	general capital sector	α	Cobb-Douglas exponent
F	fossil energy sector	τ	tax on fossil energy investments
δ	capital discount rate	$\mathcal{T}$	revenue of the fossil energy tax
R	renewable energy sector	c	additional cost of renewable energy
G	greenhouse gas level	ϕ	breakdown fraction of greenhouse gas
v	scale of climatic damage	ε	environmental tax on income
β	scaling factor	$\mathcal{E}$	fund financed by the environmental tax
t	time step		

factors that were not essential to the current study were removed, in particular elements relating to labor, technical details of global warming, and resource constraints. The reason is that our model aims to be illustrative rather than to accurately replicate reality. Moreover, simplification here allows for additional complexity in the module describing the evolution of strategies.

A fundamental difference between the present evolutionary agent-based approach and the general-equilibrium approach of Nordhaus is that here agents do not make perfectly rational decisions that are based on perfect knowledge. Instead, agents evolve their strategies through random mutation and selective imitation in a social network. Moreover, while here agents are homogeneous in terms of production functions, initial strategies and initial income, they are heterogeneous in their placement in the social network and the information they receive, and their strategies and income quickly diverge.

The numerical simulations are based on a discrete synchronous time model where the economy and strategy of each agent are updated in parallel at fixed time intervals. A time step is divided into two separate update operations: 1) updating the economy: each agent invests its income according to its own investment strategy and the individual returns are calculated by standard growth and production functions; 2) updating the strategies: all agents compare their strategies and those agent that decide to imitate change their strategies simultaneously. Each policy is evaluated over a period of 400 time steps, simulating 400 quarters or 100 years, a period that is sufficiently long to study the long-term effects of a policy on climate and welfare. As no significant financial market requires a publicly traded company to publish financial results more than 4 times a year, we consider a quarter to be the limit of feasibility to account for growth and to review an economic strategy. Given habitual behavior and organizational routines (Nelson and Winter, 1982), most economic agents will in fact review their strategy less often.

2.2. **Strategies, investment, and production.** All parameters of the economic model are summarized in Table 1. Our basic model of energy investment consists of three investment sectors: general capital K, fossil energy F, and renewable energy R. Here, the capital

accumulation in an energy sector includes technology, infrastructure, and licenses for production, distribution, and consumption of a particular form of energy. Let $Y_a(t)$ be the income of agent a at time t. Formally, the investment strategy $s_a(t)$ of an agent can be defined as a three dimensional vector

$$(1) \qquad s_a(t) = [0,1]^3, \qquad \sum_{i \in \{K,F,R\}} s_{i,a}(t) = 1.$$

The non-negative partial strategy $s_{i,a}(t)$ determines the fraction of the previous period's income $Y_a(t-1)$ that agent a invests in sector i at time t. All partial strategies are constrained to add up to one. The set of all possible investment strategies is a two dimensional simplex (i.e., a triangular surface) embedded in a three dimensional Euclidean space.

Invested capital is non-malleable: once invested it can not be transferred between sectors. The accumulation of capital in each sector depends on individual investment and the global deprecation rate δ, which is constant and equal for all sectors and all agents. In the case of fossil energy the investment can be reduced by a regulatory tax τ on investments in the fossil energy sector. This tax is defined as a fraction of fossil energy investments before taxes, so that a tax of $\tau = 100\%$ doubles the cost of all expenditures on production, distribution and consumption of fossil energy. In this way, if an agent's total spending on fossil energy is $x = Y_a(t-1)s_{F,a}(t)$, then an amount of $\frac{x}{1+\tau}$ is indeed invested, while the remaining $\frac{x\,\tau}{1+\tau}$ is paid as a tax. The revenue

$$(2) \qquad \mathcal{T}(t) = \frac{\tau}{1+\tau} \sum_a Y_a(t-1)s_{F,a}(t)$$

of this tax is recycled and distributed evenly among all agents.

To model a competitive disadvantage for renewable energy—for example through a higher cost of technology, production, or storage—we introduce an additional cost c for renewable energy, representing the difference between the costs of renewable and fossil energy. In analogy to the fossil energy tax τ, we express this additional cost in percent of the unit cost of fossil energy before taxes, i.e., renewable energy is twice as expensive as fossil energy before taxes when $c = 100\%$. The difference equations for non-aggregate growth per sector are then

$$(3) \qquad \Delta K_a(t) = Y_a(t-1)s_{K,a}(t) - \delta K_a(t-1)$$

$$(4) \qquad \Delta F_a(t) = \frac{Y_a(t-1)s_{F,a}(t)}{1+\tau} - \delta F_a(t-1)$$

$$(5) \qquad \Delta R_a(t) = \frac{Y_a(t-1)s_{R,a}(t)}{1+c} - \delta R_a(t-1).$$

We proceed by first calculating the income of agent a as if there had been no global warming, and then by accounting for global warming. We calculate $Q_a(t)$ from the returns of the individual capital sectors by a Cobb-Douglas type production function with constant

returns to scale and constant elasticity of substitution,[1]

$$(6) \qquad Q_a(t) = \beta \big(K_a(t)\big)^\alpha \big(F_a(t) + R_a(t)\big)^{1-\alpha},$$

where β is a scaling factor. In this production function fossil energy and renewable energy are assumed to be perfect substitutes: one can completely replace the other. General capital and combined energy are assumed to be imperfect substitutes. Production requires both types of input, and only a specific combination will lead to a high production level.

Global warming is commonly defined as the increase of global mean temperature above the pre-industrial mean, due to an increased level of atmospheric greenhouse gases $G(t)$. The dynamics of the greenhouse gas effect includes many local and global subsystems, resulting in complex and chaotic dynamics that allow for a range of possible climate scenarios (e.g., Stainforth et al., 2005). Here we just include a simple feedback loop that captures one of the main characteristics of greenhouse gas induced economic damage: a long delay between action and reaction that spans several decades. We do so by modeling the level of atmospheric greenhouse gases as a result of only two factors: cumulative fossil energy consumption by economic agents, which we assume to be equal to the total amount of capital accumulated in the fossil energy sector, and a natural breakdown fraction ϕ,

$$(7) \qquad \Delta G(t) \; = \; \sum_a F_a(t) - \phi \, G(t-1).$$

We assume that renewable energy does not contribute to global warming. We further pose the relationship between economic damage, global warming, and economic damage to be linear, scaled by a factor v. The net income $Y_a(t)$ of an agent a can then be calculated as

$$(8) \qquad Y_a(t) \; = \; Q_a(t) \, [1 - vG(t)] + \frac{\mathcal{T}(t-1)}{n}.$$

where $\mathcal{T}(t-1)$ are the revenues from the regulatory tax τ, distributed with one time step delay among the n agents of the population. The growth rate $\gamma_a(t)$ of the income of agent a is

$$(9) \qquad \gamma_a(t) = \frac{Y_a(t)}{Y_a(t-1)} - 1.$$

2.3. **The social network.** To model which agent can imitate which other agent we arrange all agents in a social network where the nodes are agents and the edges are communication links. We use a generic class of social networks that has been well studied and validated in network theory, namely small world networks that have a scale-free degree distribution generated by a stochastic growth process with preferential attachment (Barabási and Albert, 1999) and that have a high clustering coefficient $\mathcal{C}$.[2]

---

[1]Instead of including the fossil energy tax τ and the additional cost c of renewable energy in the growth functions, they might be incorporated in the production function,

$$Q_a(t) = \beta \, (K_a(t))^\alpha \, \left(\frac{F_a(t)}{1+\tau} + \frac{R_a(t)}{1+c} \right)^{1-\alpha}.$$

[2]In their seminal paper Watts and Strogatz (1998) define the clustering coefficient $\mathcal{C}_i$ of a node i as the number of all direct links between the immediate neighbors of i divided by the maximum number of links that could possibly exist between them. They define the clustering coefficient $\mathcal{C}$ of the entire network as the average clustering coefficient of the nodes of the network.

Before the start of each simulation we use a stochastic process to generate a new bi-directional network. The process assigns each agent a a set of peers N_a that does not change during the course of the simulation. If agent a is a peer of agent b, then a will consider the income growth rate and the investment strategy of b when choosing an agent for imitation, while b will consider the income growth rate and the investment strategy of a. On the other hand, if a and b are not peers, they will not consider each other for the purpose of imitation. The generating process starts from a circular network where each agent has two neighbors—i.e., average connectivity k =2—and iteratively adds new edges to the network until the desired average connectivity k is reached. The agents for the next new edge are chosen at random with a probability that is proportional to their connectivity (hence the term 'preferential attachment') and their proximity in the network, i.e., the inverse of the minimum number of links to traverse from one agent to the other.

The random way in which the network is created guarantees that the average distance between any two agents is very short, significantly shorter for example than in a regular grid. The preferential attachment leads to a very skewed distribution of peers per agent, with some agents having several times the median connectivity. These well connected agents act as information hubs and dominate the flow of information. A high clustering coefficient implies that if two agents are peers of the same agent, the probability that they are also peers of each other is significantly higher than the probability that two randomly chosen agents are peers. This leads to the emergence of blocks within the social network that exhibit a high level of local interconnectivity, like for example the European Union in the case of independent nations.

2.4. **Evolution of strategies.** From the point of view of evolutionary modeling, agents and investment strategies are not the same: an agent carries or maintains a strategy, but it can change its strategy and we still consider it to be the same agent (Nowak, 2006). Because every agent has exactly one strategy at a time, the number of active strategies is the same as the number of agents.

At each time step an agent may select one of its peers in the social network and imitate its strategy. If that happens, the strategy of the imitating agent changes, while the strategy of the imitated agent does not. The choice of which agent to imitate is based on relative welfare as indicated by the current growth rate of income. Note that the relation between income $Y_a(t)$ and growth $\gamma_a(t)$ is

$$(10) \qquad Y_a(t) = Y_a(0) \prod_{i=1}^{t} \left[\gamma_a(i) + 1 \right].$$

The imitating agent always selects the peer with the highest current income growth rate. Only if an agent has no peer with an income growth rate higher than itself, the agent does not revise its strategy.

If imitation were the only mechanism by which agents change their strategies, the strategies of agents that form a connected network must converge on a strategy that was present during the initial setup. However, real imitation is never without errors. Errors are called mutations in evolutionary theory. They are fundamental to an evolutionary process because they create and maintain the diversity on which selection can work. In this model we implement mutation by adding some Gaussian noise to the imitation process. That is, when an agent imitates a strategy, it adds some random noise drawn from a Gaussian distribution with zero mean to each partial strategy. This causes small mutations along each partial strategy to be more likely than large ones. The exact formula by which agent a imitates

and then mutates the strategy of agent b is

$$(11) \qquad\qquad s_a(t) = s_b(t-1) + \mathcal{N}(0, \sigma),$$

where $\mathcal{N}(0, \sigma)$ denotes a normally distributed three dimensional random vector with zero mean and standard deviation σ per dimension. Because partial investment strategies have to sum to one, we have to enforce $\mathcal{N}(0, \sigma) = 0$, for example by orthogonal projection on the simplex, resulting in the loss of one degree of freedom. The error term is further constraint to leave all partial strategies positive. Needless to say that we do not imply that our boundedly rational agents engage consciously in such mathematical exercises. Subjectively they merely allocate their income such that none is left.

In order to measure the impact of an individual agent on the evolution of strategies at the population level, we need to introduce the concept of fitness. In analogy with biology, where fitness usually measures an individual's capability to reproduce, we define the fitness of an economic agent as the frequency with which it is imitated. In the model that has been presented so far, the frequency with which agent a is imitated is fully determined by the income growth rate of a and its first and second degree neighbors. Further degrees do not matter. First degree neighbors are relevant because only direct neighbors consider a for imitation. Second degree neighbors are relevant as they are the agents that a competes with. Agent a will only be imitated by agent b if a has a higher income growth rate than b and all other neighbors of b (who are second degree neighbors of a).

This functional relationship can be expressed by a fitness function. Let $\{N_b \cup b\}$ be the set consisting of agent b and its peers, i.e., those agents with which agent b compares its income growth rate. Let $\gamma_{N_b \cup b}^{max}(t)$ be the income growth rate of the fastest growing agent in this set at time t,

$$(12) \qquad\qquad \gamma_{N_b \cup b}^{max}(t) = \operatorname*{argmax}_{c \in \{N_b \cup b\}} \gamma_c(t).$$

Then the fitness $\mathcal{F}_a(t)$ of agent a at time t is

$$(13) \qquad\qquad \mathcal{F}_a(t) = \sum_{b \in N_a} \begin{cases} 1 & \text{if } \gamma_a(t) = \gamma_{N_b \cup b}^{max}(t), \\ 0 & \text{otherwise.} \end{cases}$$

Or, in set notation:

$$(14) \qquad\qquad \mathcal{F}_a(t) = \left| \{b | b \in N_a \wedge \gamma_a(t) = \gamma_{N_b \cup b}^{max}(t)\} \right|.$$

In this function the fitness of an agent is bounded by the number of its neighbors. An agent a_1 who has just one neighbor and has the highest income growth rate among the neighbors of that neighbor has a fitness of one, whereas an agent a_{10} who has ten other neighbors and whose income growth rate is highest among the neighbors of just two of them has a fitness of two, even if in absolute terms a_1 has a much higher income growth rate than a_{10}. We see that the principal factors that determine the fitness of an agent are relative welfare as indicated by the current growth rate of income as well as the number of agents it communicates with. This gives us two different means by which a policy can regulate the evolution of economic strategies: either by changing the income growth rate of some agents, depending on the desirability of their current strategies, or by changing their connectivity in the social network, again depending on the desirability of their current strategies.

2.5. **Policy goals and formulation.** The goal of the policies that are being studied here is to let the economic agents reach a high social welfare. Assuming that fossil fuel consumption has a negative economic impact because of the associated global warming, a successful policy has to reduce consumption in fossil fuels but without considerably reducing social welfare, such that the social costs of implementing the policy do not outweigh the social benefits from a reduction in global warming.

We will study three policies, starting with a tax τ on fossil energy investments. This is the first best policy under traditionally assumed conditions (rational agents, perfect markets), and we study it here in the context of imperfect information and bounded rationality. It is a regulatory and not a revenue raising tax and is defined as a fixed percentage on all investments in fossil energy, cf. equation 2, 4, and 8. We compare this standard policy with two novel policies that take advantage of the evolutionary process by increasing the fitness of those agents that invest a larger fraction of their income in renewable energy. These policies increase the fitness of an agent either by increasing its income growth rate, or by increasing its visibility in the social network. The rationale is that, if we increase the fitness of agents that use certain strategies, these strategies will be employed more frequently.

Under the first policy, agents pay a tax that is proportional to their investment in fossil fuel. This tax makes investment in fossil energy economically less attractive. However, since the incentive not to comply with this policy is also proportional to their investment in fossil fuel, the effect of this policy depends much on the existence of a central authority that can enforce it. The second policy studied here, *prizes*, increases the fitness of agents that invest a larger fraction of their income in renewable energy by awarding them a monetary prize, financed by a global tax that is payed by all agents. That is, it is not important who pays the tax, as long as someone pays it, for example those agents that are most affected by global warming. This does not entirely solve the problem of compliance, but makes it less acute. The third policy, *advertisement*, increases the fitness of agents that invest a higher fraction of their income in renewable energy by increasing their social visibility, i.e. their connectivity in the social network. No compliance is required.

The *prizes* policy gives a monetary prize to those agents who invest the largest fraction of income in renewable energy, increasing their relative welfare, and with that their fitness. This prize is financed by an environmental tax ε on production $Q_a(t)$. Since this is a revenue raising tax to finance the policy and not a regulatory tax that depends on individual behavior, it has the same level for each agent. Let $\mathcal{E}(t)$ be the size of the environmental fund at time t:

$$(15) \qquad \mathcal{E}(t) = \sum_a Q_a(t) \left[1 - vG(t)\right] \varepsilon.$$

At each time step, the q agents that invest the highest fraction of their income in renewable energy are each awarded an equal share $\mathcal{E}(t-1)/q$, such that under the *prizes* policy the net income becomes

$$(16) \qquad Y_a(t) = Q_a(t) \left[1 - vG(t)\right] \left[1 - \varepsilon\right] + \begin{cases} \mathcal{E}(t-1)/q & \text{if } a \text{ is awarded a prize,} \\ 0 & \text{otherwise.} \end{cases}$$

To give an example: if the income tax ε is 1%, and 10 out of 200 agents are selected to receive a prize, then under the assumption that their income does not deviate significantly from the average income, it is raised by about 20%. If the majority of agents receive a prize, the tax to finance the prizes is in effect a selective punishment of those agents that invest relatively much in fossil fuels.

TABLE 2. Free policy parameters

Fossil energy tax	τ	tax on fossil energy investments
Prizes	q	number of agents that receive a prize
	ε	tax on income to finance the prize
Advertisement	q	number of agents that are advertised
	p	probability that an agent is reached by advertisement (the simulations use a fixed value of $p = .25$)

TABLE 3. Calibrated parameter values of the economic model

k	network connectivity	10	$\mathcal{C}$	clustering coefficient	.66
σ	mutation variance	.02	α	Cobb-Douglas exponent	.9
δ	capital discount rate	.01	ϕ	breakdown of greenhouse gases	.01
β	scaling factor	.021	v	scale of climatic damage	.00007

The *advertisement* policy increases the social visibility of those agents that invest the largest fraction of their income in the renewable energy sector, increasing the number of agents that consider the advertised agents when deciding whom to imitate. At each time step the q agents that invest the largest fraction of their income in renewables are selected to be advertised. The advertised agents are temporarily added to the group of neighbors of some other agents, so that these other agents consider the advertised agents when deciding whom to imitate. Advertisement does not oblige an agent to consider an advertised agent. Instead, its success rate depends on the resources invested in the campaign. For simplicity we assume that whether agent a considers the advertised agent b for imitation is an independent random event for each a, b and t and has probability p. We ignore the cost of advertisement and assume a success rate of just $p = .25$.

To give an example, let the average number of neighbors per agent before advertisement be $k = 10$ and let $q = 8$ agents be selected for advertisement. On average, each agent can now choose between $k + q * p = 10 + 8 * .25 = 12$ neighbors when deciding whom to imitate. If an agent imitates, chances are one in six that it imitates the strategy of an advertised agent, provided that the income of the advertised agents does not deviate significantly from that of the other agents. The free parameters of each policy are listed in Table 2.

2.6. **Model calibration.** The free parameters of the economic model are calibrated such that global warming has a significant negative welfare effect, emphasizing the need for policies. The calibrated values of all free economic parameters are summarized in Table 3.

A fixed number 200 of agents is used in all simulations; this is approximately the number of independent states and a rough approximation of the number of agents with an independent energy policy. The quarterly capital discount rate is $\delta = .01$. The exponent of general capital in the production function is $\alpha = .9$, and the exponent of the combined energy sector is $1 - \alpha = .1$. In this way income is highest when 90% of an agent's capital is in the general sector and 10% in the two energy sectors. The scaling factor β of the production function is calibrated such that the calibrated economic model without climate
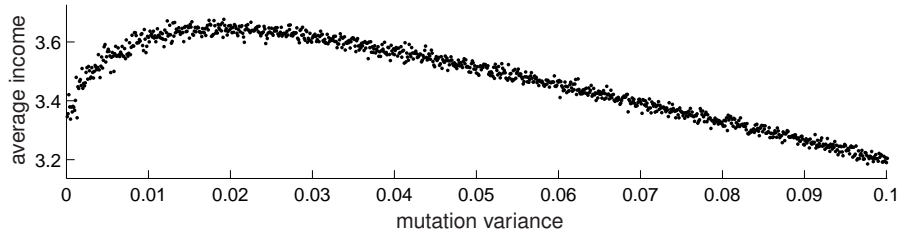
FIGURE 1. Effect of the mutation variance on economic performance

damage has an economic growth rate of about 2% per annum. The breakdown fraction ϕ of greenhouse gas and the sensitivity v to global warming are calibrated such that without any climate policy the greenhouse gas emissions reduce the per annum growth rate by an order of magnitude over the 100 years of the simulation, consistent with the studies reported in Section 1.

The mutation variance σ is the only free parameter that regulates the evolutionary mechanism. Small values of σ slow down the discovery of a good strategy. Large values prevent convergence. A good value of σ lies somewhere in between. Figure 1 shows how the average income of the agents depends on σ. The x-axis shows different values for σ. The y-axis shows the average income that a population of 200 agents realize after 400 time steps. Each measurement point in the graph is averaged over 100 simulations. The initial strategy of each agent is chosen at random. There are no taxes, global warming has no effect, and the additional cost of renewable energy is $c = 100\%$. Under these conditions the optimal strategy that maximizes the income growth rate of an individual agent is $\langle s_K, s_F, s_R \rangle = \langle .9, .1, 0 \rangle$. The graph shows that average income is maximized for a value of approximately $\sigma \approx .02$, and for this reason we use a value of $\sigma = .02$ in the remainder of this study.

For the social network we use an average connectivity of $k = 10$. In a population of 200 agents this value results in a highly connected network—the average distance between any two agents in the network is 2.7—while maintaining the overall qualities of a complex social network. Figure 2 shows some key statistics collected from 10,000 networks of 200 agents that were created by the stochastic growth process using these values: a normalized histogram of the clustering coefficient $\mathcal{C}$ of each network (average $\mathcal{C}$ is .66), a normalized histogram of the distance between any two agents in each network, the probability density function (PDF), and the cumulative density function (CDF) of the number of neighbors per agent in each network. Note the relatively high probability of having 20 or more neighbors when the average connectivity is 10 neighbors. Such significant numbers of highly connected agents do not exist in regular grid networks or random networks of the Erdős-Rényi type, yet their existence in real social networks is well established (Albert and Barabási, 2002). They generally act as information or transportation hubs and accelerate the dissemination of goods, viruses and ideas.

In order to avoid any dependency of the simulation results on initial conditions, the numeric simulations are divided into an initialization phase and a main experimental phase.
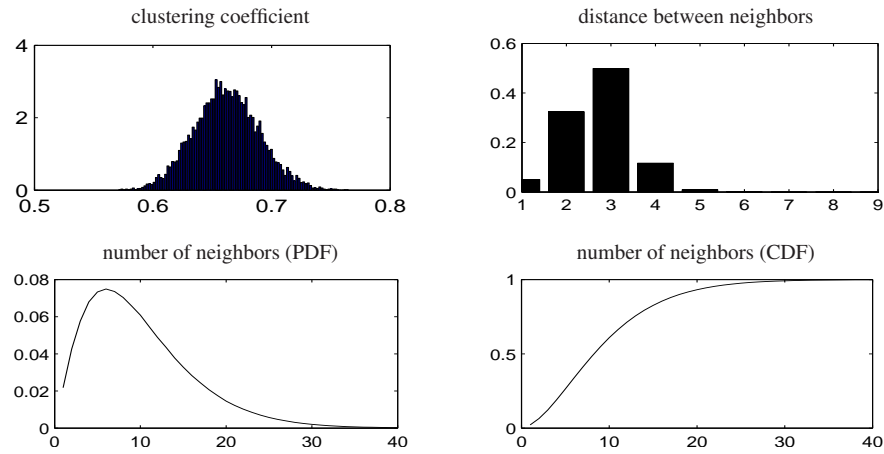
FIGURE 2. Network statistics. The two upper graphs are normalized histograms.

During the initialization phase certain parameters of the evolutionary economy will converge regardless of the initial conditions, contributing to the general validity of the numerical results. The initialization effect is visualized in Figure 3, which shows how the average investment strategy converges on an equilibrium. 800 time steps are shown. Initial strategies are chosen at random from the two dimensional simplex, and so at $t = 1$ the average strategy is $\langle s_K, s_F, s_R \rangle = \langle 1/3, 1/3, 1/3 \rangle$. The average strategy at $t = 800$ is $\langle .9, .09, .01 \rangle$. Results are averaged over 10,000 simulations.
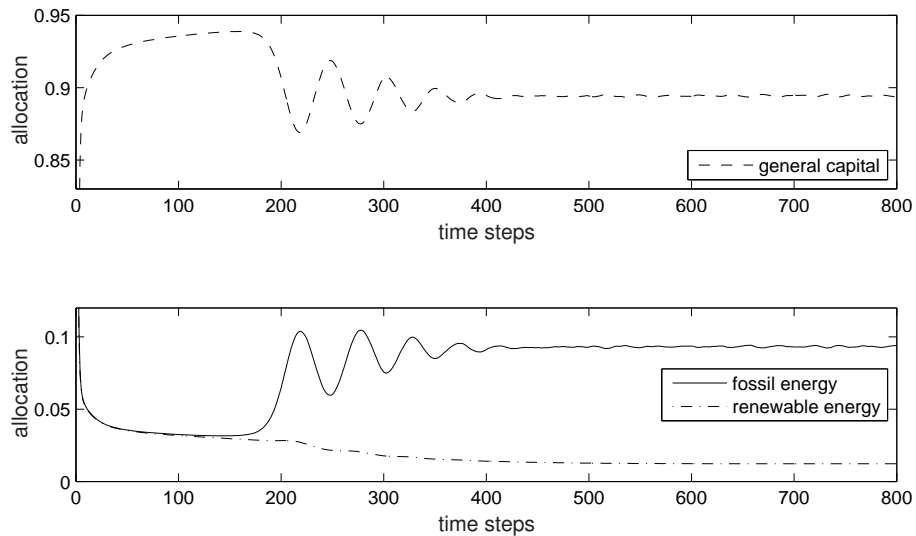
FIGURE 3. Convergence of the average investment strategy

Note that while it takes the agents only about a dozen time steps to learn to invest some 90% of their investments in general capital, they need about 200 time steps to become sufficiently sensitive to the difference in cost between the two energy sectors and to differentiate their energy investments. From $t = 200$ to $t = 400$ the convergence on the final strategy can be seen to follow a damped oscillation pattern. The full effect of a policy can only be established if it is introduced after the system without policy has reached equilibrium. It takes 400 time steps for the system without a policy to converge, and so we base the numeric evaluation of climate policies on simulations that consist of an initialization phase of 400 time steps during which no policies are applied, followed by a main experimental phase of 400 time steps during which policies are applied and evaluated. In particular, the tax τ on fossil energy investments and the environmental tax ε are always zero up until $t = 400$. Only from $t = 401$ they take the value assigned to them by the respective policy.

3. The evolutionary dynamics

3.1. Derivation of the growth function. An important prerequisite for regulating an evolutionary system is to understand its dynamics. Here we are primarily interested in what strategy the agents will converge on. With regard to global warming we are further interested in whether the evolutionary agents can converge on a globally optimal strategy, rather than on individually optimal strategies.

While the fitness function describes how the imitation of a strategy (the genotype) depends on welfare as indicated by the income growth rate (the phenotype), the *growth function* describes how a strategy determines the income growth rate of the agent that carries it. The growth function calculates the economic utility of a strategy as the equilibrium growth rate to which the income growth rate of an agent converges if it holds on to that particular strategy. Derivation of the growth function is essential for an understanding of the evolutionary dynamics. We will base it on an analysis of the ratio of sector specific capital to income.

To start with, equation 4 and 5 can be combined to express the difference equation of a combined energy sector $E = F + R$,

$$(17) \qquad \Delta E_a(t) \ = \ Y_a(t-1) \ \left(\frac{s_{F,a}(t)}{1+\tau} + \frac{s_{R,a}(t)}{1+c} \right) \ - \ \delta E_a(t-1),$$

where the combined energy investment strategy of an agent is $s_{F,a}(t) + s_{R,a}(t) = 1 - s_{K,a}(t)$. Let $r_a(t)$ be the fraction of $1 - s_{K,a}(t)$ that is invested in fossil energy, and $1 - r_a(t)$ the fraction that is invested in renewable energy,

$$(18) \qquad r_a(t) = \frac{s_{F,a}(t)}{1 - s_{K,a}(t)}.$$

This enables us to rewrite equation 17 as

$$(19) \qquad \Delta E_a(t) \ = \ Y_a(t-1) \ [1 - s_{K,a}(t)] \, f(r_a, t) \ - \ \delta E_a(t-1),$$

where $f(r_a, t)$ stands for

$$(20) \qquad f(r_a, t) \ = \ \frac{r_a(t)}{1+\tau} + \frac{1 - r_a(t)}{1+c}.$$

We collapse the scaling factor and the economic effect of global warming into a single factor $\zeta(t)$,

$$(21) \qquad \zeta(t) = \beta \left[1 - vG(t) \right].$$

Next we combine the calculation of income (equation 8) with the production function (equation 6) and simplify it by ignoring the additive term $\frac{\mathcal{T}(t-1)}{n}$, which is identical for all agents,

$$(22) \qquad Y_a(t) = \zeta(t) \, K_a(t)^\alpha \, E_a(t)^{1-\alpha}.$$

We now use equation 3 and 9 to calculate the difference equation of the ratio of general capital to income as

$$(23) \qquad \begin{aligned} \frac{K_a(t)}{Y_a(t)} &= \frac{Y_a(t-1)\, s_{K,a}(t) \ + \ (1-\delta)K_a(t-1)}{(\gamma_a(t)+1)\, Y_a(t-1)} \\[2ex] &= \frac{s_{K,a}(t)}{\gamma_a(t)+1} + \frac{1-\delta}{\gamma_a(t)+1} \, \frac{K_a(t-1)}{Y_a(t-1)}. \end{aligned}$$

This dynamic equation is of the form

$$(24) \qquad x(t) = a + bx(t-1),$$

which under the condition $0 \leq b < 1$ converges monotonically to its unique stable equilibrium at

$$(25) \qquad \lim_{t\to\infty} x(t) = a/(1-b).$$

In a model without global warming this condition is normally fulfilled: investment is always non-negative and sector specific capital cannot decrease faster than δ. Excessive economic damage caused by global warming, $vG(t)$, does theoretically allow for $\gamma_a \leq -\delta$. However, the social and political ramifications of such a catastrophic decline go beyond the scope of this model. Hence, with the restriction that this model only covers the case $\gamma_a > -\delta$, and considering that $0 < \delta \leq 1$, we have the required constraint for convergence

$$(26) \qquad 0 \leq \frac{1-\delta}{\gamma_a(t)+1} < 1.$$

We conclude that the ratio of general capital to income converges to

$$(27) \qquad \begin{aligned} \lim_{t\to\infty} \frac{K_a(t)}{Y_a(t)} &= \lim_{t\to\infty} \frac{s_{K,a}(t)}{\gamma_a(t)+1} \Big/ \left(1 - \frac{1-\delta}{\gamma_a(t)+1}\right) \\ &= \lim_{t\to\infty} \frac{s_{K,a}(t)}{\gamma_a(t)+\delta}. \end{aligned}$$

Equation 27 describes a unique stable equilibrium to which the ratio of general capital to income converges monotonically. A similar result can be obtained for the energy sector:

$$(28) \qquad \lim_{t\to\infty} \frac{E_a(t)}{Y_a(t)} = \lim_{t\to\infty} \frac{[1 - s_{K,a}(t)]\, f(r_a, t)}{\gamma_a(t)+\delta}.$$

Ignoring the limit notation we combine equation 27 and 28 with equation 22 to calculate income at equilibrium as

$$(29) \qquad \begin{aligned} Y_a(t) &= \zeta(t) \left(\frac{Y_a(t-1)\, s_{K,a}(t)}{\gamma_a(t)+\delta}\right)^\alpha \left(\frac{Y_a(t-1)\,[1-s_{K,a}(t)]\, f(r_a,t)}{\gamma_a(t)+\delta}\right)^{1-\alpha} \\ &= \zeta(t)\, \frac{Y_a(t-1)}{\gamma_a(t)+\delta}\, s_{K,a}(t)^\alpha\, [1-s_{K,a}(t)]^{1-\alpha}\, f(r_a,t)^{1-\alpha}. \end{aligned}$$

Solving for $\gamma_a(t)$ yields the growth function

$$(30) \qquad \gamma_a(t) = \zeta(t)\, s_{K,a}(t)^\alpha\, [1-s_{K,a}(t)]^{1-\alpha}\, f(r_a,t)^{1-\alpha} \; - \; \delta.$$

3.2. **Convergence behavior.** We can now address the question whether evolutionary agents can be expected to converge on the globally rather than on the individually optimal strategy. In equation 27 and 28 neither the rate of convergence nor the equilibrium itself depend on the value of $\zeta(t)$. In equation 30 we find that $\zeta(t)$ is a multiplicative factor that does not change the relative order of the equilibrium growth rate of individual strategies. Since the fitness of an agent depends on the order of income growth rates, the fitness function is invariant under such a monotonous transformation. In other words, $\zeta(t)$ does not change the likelihood of a particular strategy to be imitated. This means that global warming has no effect on the evolutionary process: agents must not be expected to show any type of behavioral response to the economic effects of global warming and are not likely to choose the globally over the individually best strategy.
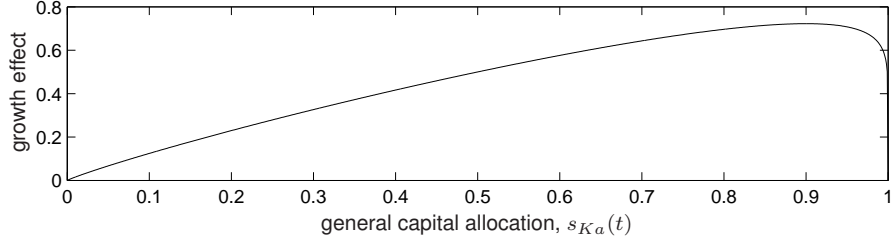
FIGURE 4. Growth effect of investment in general capital. The production coefficient is $\alpha = .9$.

To answer the question of which strategy the agents will converge on, the growth function of equation 30 can be decomposed into a term that describes the effect of income allocation to general capital on growth, and a term that describes the growth effect of the allocation of the remaining income over the two energy sectors. The dependency of the equilibrium growth rate on the general capital allocation as seen in equation 30 is given by the term

$$(31) \qquad s_{K,a}(t)^{\alpha} \left[1 - s_{K,a}(t)\right]^{1-\alpha},$$

which depends exclusively on the constant production coefficient α. This term is maximized for $s_{Ka}(t) = \alpha$, which implies that the optimal allocation to the combined energy sector is $s_{Fa}(t) + s_{Ra}(t) = 1 - \alpha$. As can be seen in Figure 4, the growth effect is a concave function of $s_{Ka}(t)$ with an extended region around the maximum that has a gradient close to zero.

The effect of $r_a(t)$ on the income growth rate is via $f(r_a, t)^{1-\alpha}$ which, from equation 20, is

$$(32) \qquad f(r_a, t)^{1-\alpha} = \left(\frac{r_a(t)}{1+\tau} + \frac{1 - r_a(t)}{1+c} \right)^{1-\alpha}.$$

Figure 5 shows that this function is flat when $\tau = c$ and otherwise concave. For $\tau > c$ the term is maximized when $r_a(t) = 0$, and for $\tau < c$ it is maximized when $r_a(t) = 1$. For a given value of $r_a(t)$ the curvature increases with $|\tau - c|$. For a given value of $|\tau - c|$ the curvature increases with the distance to the maximum. Unlike term 31, the curvature at the maximum is not zero. Maximizing this type of growth function poses no challenge to a (collective) learning mechanism. It has a single global optimum, no local optima, and a distinct slope that increases with distance to the optimum. Even the simplest of hill climbing algorithms can find and approach this optimum. Learning mechanisms will differ mostly in the speed of convergence.

Regarding the speed with which our evolutionary agents converge on the optimum strategy, we must bear in mind that evolutionary agents will only converge on the individually optimum strategy if there is sufficient selection pressure. Figure 3 shows that the speed of convergence gradually decreases as the optimum strategy is approached. The previous discussion has shown that the slope of the growth function monotonically decreases as the maximum is approached. This apparent correlation between the speed of convergence and the slope of the growth function can be explained from the fact that the actual income
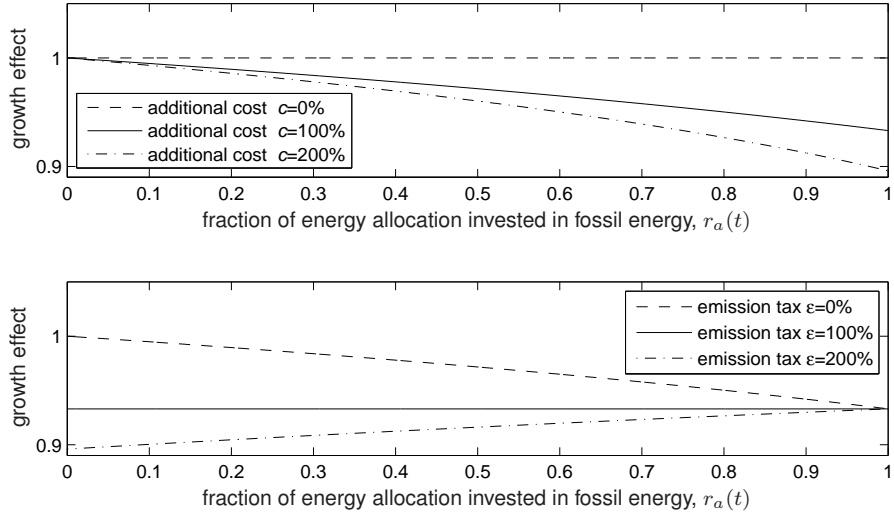
FIGURE 5. Growth effect of different investment allocations. The production coefficient is $\alpha = .9$ and the total energy allocation is $1 - s_{K,a}(t) = .1$. In the upper graph the tax on fossil energy investments is $\tau = 0$. In the lower graph the additional cost of renewable energy is $c = 100\%$.

growth rate of an agent only approximates the equilibrium growth rate of its strategy. Due to this inaccuracy, the selective advantage of an investment strategy over a variant with lower equilibrium growth rate diminishes as the difference in equilibrium growth rates decreases. So as the slope of the growth function decreases around the optimum, the selection pressure among variants decreases as well, with the important consequence that the evolutionary economy potentially never converges and never reaches equilibrium.

4. POLICY ANALYSIS

4.1. **Experimental setup.** We use numerical simulations to determine how sensitive the three policies of Section 2.5 are to a particular choice of values for their free parameters (cf. Table 2), and how sensitive they are to a particular choice of value for the cost of renewable energy. We measure this sensitivity with regard to how effective each policy is in regulating the economic behavior in an evolutionary economy, which in this model means to reduce global warming and increase social welfare.

To reduce the variance of the simulation results, we replicate each simulation 1,000 times for each tested level of free parameters and cost of renewable energy. In order to obtain results that are valid for the general class of scale-free social networks with a high clustering coefficient, each replication uses a different random instance of such a network, so that the results are valid for our general class of social networks but do not depend on a specific choice of network. Also, at the start of each replication the agents are initialized with random strategies that converge during the initialization phase of 400 time steps. During the following 400 time steps, the main experimental phase, the policy is applied. We
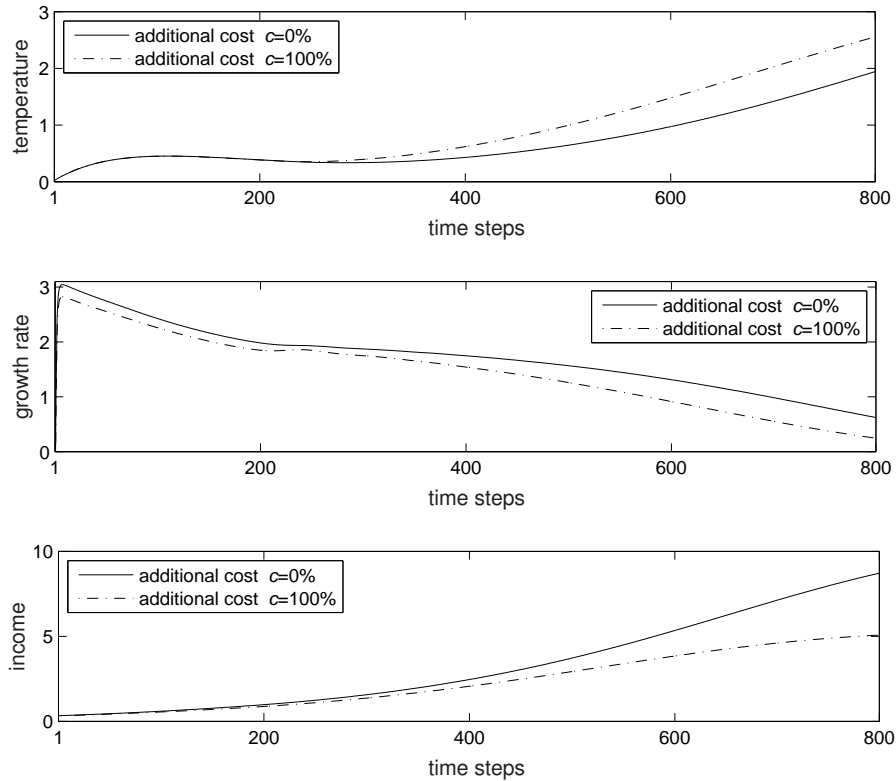
FIGURE 6. Evolution of the calibrated economy

report the average value at time $t = 800$ of three key statistics: global temperature, average income, and average income growth rate. We also report the average energy allocation at $t = 800$.

Figure 6 shows the evolution of the three key statistics when no policy is applied, for an additional cost of renewable energy $c = 0\%$ and $c = 100\%$. These are the two systems against which the policies are evaluated. For each policy and for each parameter scan, the graph will include the same statistic for a system without policy.

4.2. **Evaluating the first best policy, a tax on fossil energy investment.** When fossil energy and renewable energy are perfect substitutes, investment in the more cost-effective energy sector generates a higher income growth rate for an investing agent. A rational agent is expected to use the strategy with the highest equilibrium growth rate and to invest exclusively in the more cost efficient energy sector, even if the difference is very small: if $\tau < c$, a rational agent invests only in fossil energy. If $\tau > c$, it invests only in renewable energy. If $\tau = c$, it is indifferent between the two energy sectors. This does not hold for evolutionary agents, which converge on a strategy only if there is sufficient selection pressure. When the cost difference between fossil and renewable energy is small, the resulting difference in the equilibrium income growth rate is also small. Since the actual

income growth rate only approximates the equilibrium growth rate to a certain degree, small difference in equilibrium growth rate are harder to detect than large ones.

Figure 7 shows the economic effect at $t = 800$ of different levels of a tax on fossil energy investment. Here the additional cost of renewable energy is $c = 100\%$. The average energy allocation is a smooth function of the cost difference of the two energy sectors, and hence of the slope of equation 32. The curves can best be described as two symmetric sigmoids that cross each other at about $\tau = 125\%$. In other words, the evolutionary agents are indifferent between the two energy sectors at a tax level of $\tau = 125\%$. For a rational agent as described above, we would observe two step functions that cross each other at the point where both energy sectors carry the same cost, i.e., $\tau = 100\%$. Figure 8 allows for a similar observation for different levels of the cost of renewable energy when the tax is $\tau = 100\%$. Here the two curves (approximate sigmoids) of the energy allocation cross each other at a tax level of $c = 80\%$. Again, for a rational agent as described above, we would observe two step functions that cross each other when the additional cost of renewable energy is equal to the tax, i.e., at $c = 100\%$.

That the observed crossover points deviate significantly from the point $\tau = c$ where rational agents are indifferent is due to a particular type of lock-in or memory effect of an evolutionary system. During the initialization phase no policy is applied and due to its selective advantage the agents converge on a strategy that allocates energy investments to fossil energy. During the main experimental phase a tax on fossil energy investment changes the selective advantage in favor of renewable energy. As the agents move towards the new optimum, the slope of the growth function decreases to the point that the selection pressure becomes insignificant. For all practical purposes, the convergence comes to a halt somewhere between the old and the new optimum.

In both Figure 7 and 8 the increase in global temperature generally reflects the investment in fossil energy, and this increase is in turn reflected in the average income growth rate and the average income. All statistics change monotonically as a function of $\tau - c$. The higher the tax on fossil energy investment, the lower the global temperature and the higher the average income growth rate and the average income.

4.3. **Evaluating the *prizes* policy.** In a model of rational expectations, a prize that is awarded only to selected agents introduces complex social dynamics that can be highly sensitive to initial conditions or, worse, intractable (Challet et al., 2005). In this evolutionary model, agents do not make choices in anticipation of a prize. Instead, they choose to imitate an agent after a prize is given, based on relative welfare. The evolutionary impact of a prize is a simple function of its effect on the relative welfare.

Figure 9 shows how the economic impact of a prize varies with the number q of rewarded individuals. Here the income tax that finances the prize is $\varepsilon = 1\%$ and the additional cost of renewable energy is $c = 100\%$. When $q = 3$ agents are rewarded, both the average investment in fossil fuel and the global temperature are minimized, and average growth and income are maximized. For values of $10 \leq q \leq 50$ the agents weakly prefer renewable energy, and for all values of $q \leq 160$ a significant improvement in income growth rate and income can be observed, compared to the system without policies. Very high values of q have no positive economic effect, and we conclude that the evolutionary system is not showing the same positive response to selective punishment as it shows to selective reward.

Figure 10 shows the policy effect for different levels of income tax, for an additional cost of renewable energy $c = 100\%$ and $q = 3$ rewarded agents. Investments in renewable

energy increase monotonically as the tax increases, and the global temperature decreases. The positive welfare effect however peaks at a tax of 3% and declines for higher tax levels.

Figure 11 shows how the policy effect varies with the cost of renewable energy, for $q = 3$ rewarded agents and an environmental tax on income $\varepsilon = 1\%$. With the chosen parameters the policy proves to be effective for an additional cost of renewable energy of up to 100%. While the global temperature and the average income growth rate are positively affected even by higher cost levels, average income approaches that of the system without policy.

4.4. **Evaluating the *advertisement* policy.** The social effect of *advertisement* can not be understood correctly without the concept of evolutionary fitness. No money is being transferred and there is no increase in information. All that is changed is the number of other agents that consider an advertised agent for imitation.

Figure 12 shows how the economic impact of *advertisement* varies with the number q of advertised agents. The additional cost of renewable energy is $c = 0\%$. A broad range of values for the number q of advertised agents proves to be effective, peaking in the region of ten to forty agents, and decreasing slowly as the maximum of $q = 200$ is approached, at which point the network is fully connected. Figure 13 shows how the effect of the advertisement policy varies with the cost of renewable energy. The number of advertised agents is $q = 10$. The policy proves to be effective only up to an additional cost of renewable energy of 1%. Beyond this point global temperature and average income approach the levels without policy, and the income growth rate becomes even lower than without policy. In other words, *advertisement* is only effective when the slope of the growth function is small (cf. equation 32) and the selection pressure to invest in the more cost efficient fossil energy sector is negligible.

5. SUMMARY AND CONCLUSIONS

An agent-based simulation of an economic process facilitates the study of climate policies under conditions that are difficult if not impossible to study in equilibrium type of models with representative and rational agents. The agent-based approach describes agents that are heterogeneous in their strategies and assets and reflect bounded rationality. This allows for the implementation and study of selective policies that treat agents differently depending on their behavior. The particular form of the agent-based model employed here is an evolutionary model of strategy formation in a social network.

The approach was applied to model investment choices by individual agents in general, fossil energy, and renewable energy capital, as part of a simple economic model with global warming feedback. Use of fossil energy is the individually optimal strategy, but causes global warming and a decline in social welfare that calls for a climate policy. As there is no central authority to enforce a climate policy, compliance is a problem. A social-evolutionary module describes selective imitation and random mutation of investment strategies. The probability that an agent is imitated depends on relative welfare and social connectivity, with relative welfare measured by the relative growth rate of individual income in the individual's (peer) network.

How an investment strategy translates into relative welfare as indicated by the income growth rate is described by a growth function. In this economic model the growth function is a concave function with a single optimum, in principle an easy optimization problem for any type of learning algorithm. However, as the strategies converge on the optimum, selection pressure decreases, the convergence process slows down, and the evolutionary economy potentially never reaches equilibrium. The growth function shows further that

global warming has no effect on relative welfare, and the evolutionary agents therefore can not be expected to choose the globally over the individually optimal strategy.

Two selective policies were formulated that take heterogeneity of the strategies and of the social connectivity of individual agents into account. They influence the evolutionary formation of strategies by increasing the probability of desirable strategies to be imitated. Numerical simulations compared both policies with that of a standard regulatory tax on fossil energy investment, measuring how effective each policy is in reducing global warming and increasing social welfare. One selective policy, *prizes*, regulates relative welfare positions and causes agents with a desirable strategy to be ranked higher by their neighbors. The other selective policy, *advertisement*, regulates social visibility so that agents with a desirable strategy are seen by more agents. With regard to effectiveness, the regulatory tax on fossil energy investment depends on the compliance of the big polluters. *Prizes* depends on the compliance of at least some agents that pay into the environmental fund, for example those agents that suffer most from global warming. *Advertisement* does not depend on enforcement.

Both *prizes* and the tax on fossil energy investment were found to be effective over a wide range of values for the additional cost of renewable energy, with a gradual decrease in effectiveness as this cost increases. Numerical evaluation of the tax on fossil energy investment has shown that due to lock-in, the tax level at which evolutionary agents are indifferent between the two energy investment sectors is significantly higher than the tax level at which their costs are equal. This can be seen to reflect a tax on a lock-in externality. *Prizes* have shown that an evolutionary system is far less responsive to selective punishment than to a prize. *Advertisement* only works well when the cost difference between the two energy sectors is very small and the selection pressure to invest in fossil energy is very low.

The evolution of economic strategies and the dynamics of global warming are far more complex than expressed here, but one may expect that selective policies have the same qualitative effect. The effect of *prizes* is similar to that of a regulatory tax on fossil energy investment, but depends less on the compliance of the big polluters. When the costs of fossil and renewable energy are nearly equal, economic fairs and conferences, scholarships, awards for outstanding contributions, and publication of informative material represent relatively inexpensive policy tools that do not depend on an enforcing authority and that can have a significant effect on combating global warming.

REFERENCES

Albert, R., Barabási, A.-L., 2002. Statistical mechanics of complex networks. Rev. Modern Phys. 74 (1), 47–97.

Barabási, A. L., Albert, R., 1999. Emergence of Scaling in Random Networks. Science 286, 509–511.

Challet, D., Marsili, M., Zhang, Y.-C., 2005. Minority Games: Interacting Agents in Financial Markets. Oxford University Press, Oxford.

Dopfer, K. (Ed.), 2005. The Evolutionary Foundation of Economics. Cambridge University Press, Cambridge, England.

Epstein, J. M., Axtell, R., 1996. Growing Artificial Societies: Social Science from the Bottom Up. MIT Press, Cambridge, MA.

Erdős, P., Rényi, A., 1959. On random graphs. Publ. Math. 6, 290–297.

Ferrer-í-Carbonell, A., 2005. Income and Well-being: An Empirical Analysis of the Comparison Income Effect. J. Public Econ. 89 (5–6), 997–1019.

Frank, R. H., 1987. Choosing the Right Pond: Human Behavior and the Quest for Status. Oxford University Press, Oxford.

Friedman, D., 1998. On economic applications of evolutionary game theory. J. Evol. Econ. 8, 15–43.

Janssen, M., Anderies, J. M., Walker, B. H., 2004. Robust strategies for managing rangelands with multiple stable attractors. J. Environ. Econ. Manag. 47, 140–162.

Janssen, M., Ostrom, E., 2006. Governing social-ecological systems. In: Tesfatsion and Judd (2006), Ch. 30, pp. 1465–1509.

Kelly, D. L., Kolstad, C. D., 1999. Integrated assessment models for climate change control. In: Folmer, H., Tietenberg, T. (Eds.), The International Yearbook of Environmental and Resource Economics 1999-2000. Edward Elgar Publishing, Cheltenham, pp. 31–98.

Levy, H., Levy, M., Solomon, S., 2000. Microscopic Simulation of Financial Markets: From Investor Behavior to Phenomena. Academic Press, New York.

Manne, A. S., 1992. Buying Greenhouse Insurance: The Economic Costs of CO2 Emission Limits. MIT Press, Cambridge, MA.

Nelson, R. R., Winter, S. G., 1982. An Evolutionary Theory of Economic Change. Harvard University Press, Cambridge, MA.

Nordhaus, W. D., 1991. To Slow or Not to Slow: The Economics of the Greenhouse Effect. Econ. J. 101 (407), 920–937.

Nordhaus, W. D., 1992. An Optimal Transition Path for Controlling Greenhouse Gases. Science 258, 1315–1319.

Nordhaus, W. D., 1994. Managing the global commons: The economics of the greenhouse effect. MIT Press, Cambridge, MA.

Nordhaus, W. D., Yang, Z., 1996. A Regional Dynamic General-Equilibrium Model of Alternative Climate-Change Strategies. Amer. Econ. Rev. 86 (4), 741–765.

Nowak, M. A., 2006. Evolutionary Dynamics: Exploring the Equations of Life. Harvard University Press, Cambridge, MA.

Peck, S. C., Teisberg, T. J., 1993. Optimal CO2 Emissions Control with Partial and Full World-wide Cooperation: An Analysis Using CETA. Electric Power Research Institute, Palo Alto, Calif.

Popp, D., 2004. ENTICE: endogenous technological change in the DICE model of global warming. J. Environ. Econ. Manag. 48, 742–768.

Stainforth, D. A., Aina, T., Christensen, C., Collins, M., Faull, N., Frame, D. J., Kettleborough, J. A., Knight, S., Martin, A., Murphy, J. M., Piani, C., Sexton, D., Smith, L. A.,

Spicer, R. A., Thorpe, A. J., Allen, M. R., 2005. Uncertainty in predictions of the climate response to rising levels of greenhouse gases. Nature 7, 1–22, 224–254.

Tesfatsion, L., Judd, K. L. (Eds.), 2006. Agent-based Computational Economics. Vol. 2 of Handb. Comput. Econ. North-Holland, Amsterdam.

Tol, R. S. J., 1995. The Climate Fund: Sensitivity, Uncertainty, and Robustness Analysis. Tech. Rep. W95/02, Institute for Environmental Studies, Vrije Universiteit, Amsterdam.

van den Bergh, J. C. J. M., 2004. Optimal climate policy is a utopia: from quantitative to qualitative cost-benefit analysis. Ecol. Econ. 48, 385–393.

Watts, D. J., Strogatz, S. H., 1998. Collective dynamics of 'small-world' networks. Nature 373, 440–442.

Wilhite, A. W., 2006. Economic activity on fixed networks. In: Tesfatsion and Judd (2006), Ch. 20, pp. 1013–1045.

Witt, U., 2008. What is specific about evolutionary economics? J. Evol. Econ. 18, 547–575.
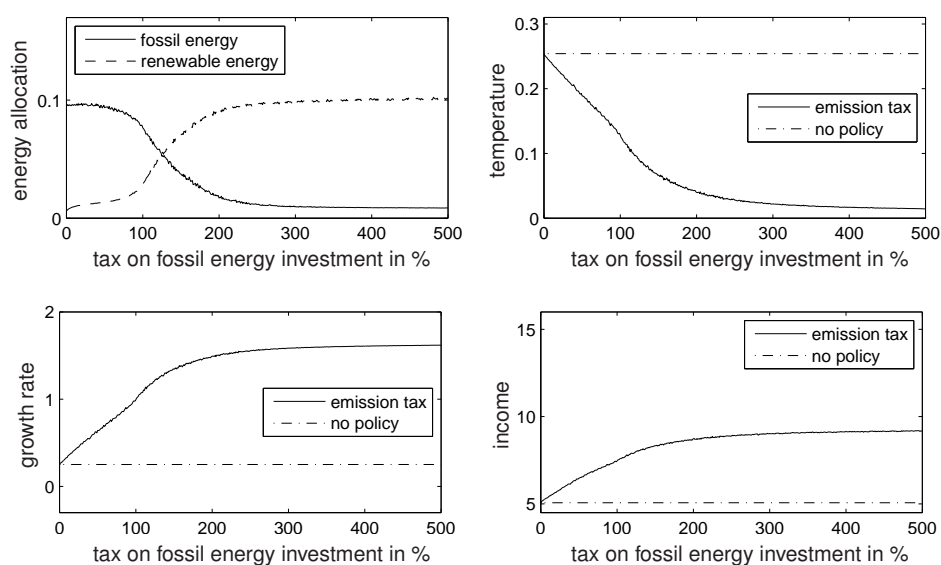
FIGURE 7. Effect of a tax on fossil energy investment for different tax levels, at $t = 800$. The additional cost of renewable energy is $c = 100\%$.
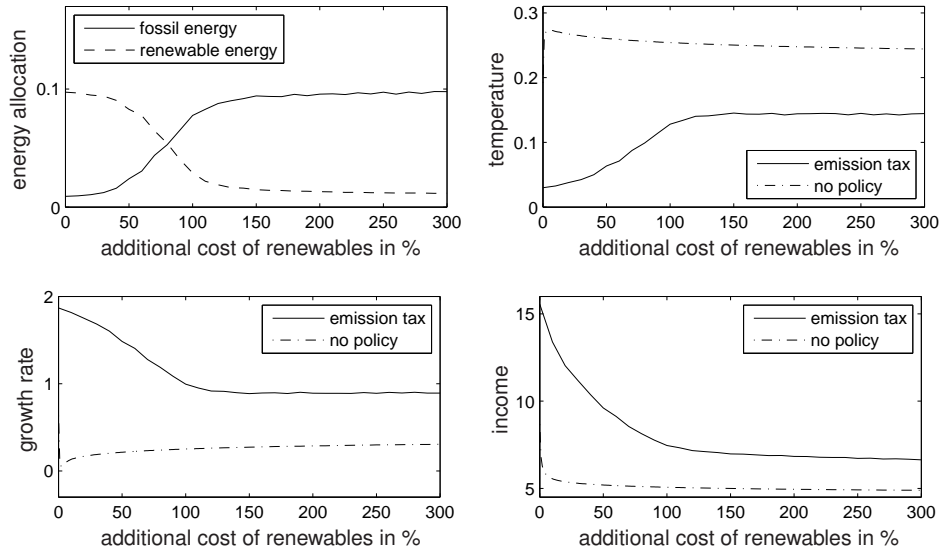
FIGURE 8. Effect of a tax on fossil energy investment for different levels of cost of renewable energy, at $t = 800$. The tax level is $\tau = 100\%$. Results are averaged over 1,000 simulations for every cost increment of 10 percent points.

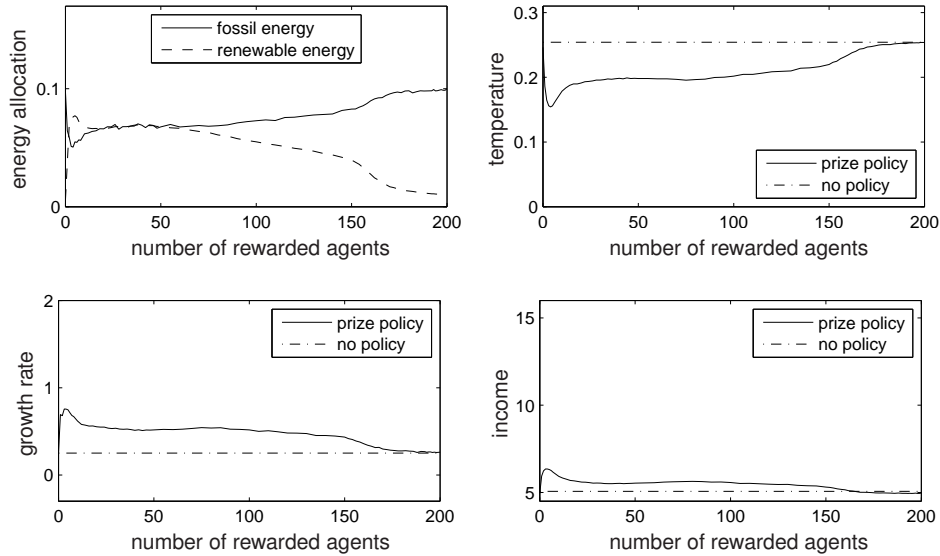

FIGURE 9. Effect of the *prizes* policy for different numbers of rewarded agents, at $t = 800$. The additional cost of renewable energy is $c = 100\%$. The environmental tax on income is $\varepsilon = 1\%$.
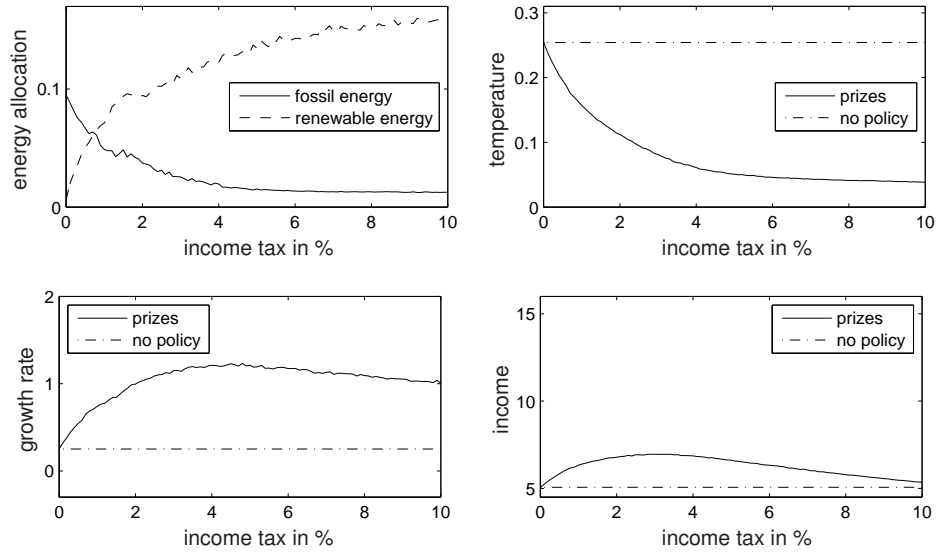
FIGURE 10. Effect of the *prizes* policy for different levels of the environmental tax on income, at $t = 800$. The additional cost of renewable energy is $c = 100\%$, the number of rewarded agents is $q = 3$.
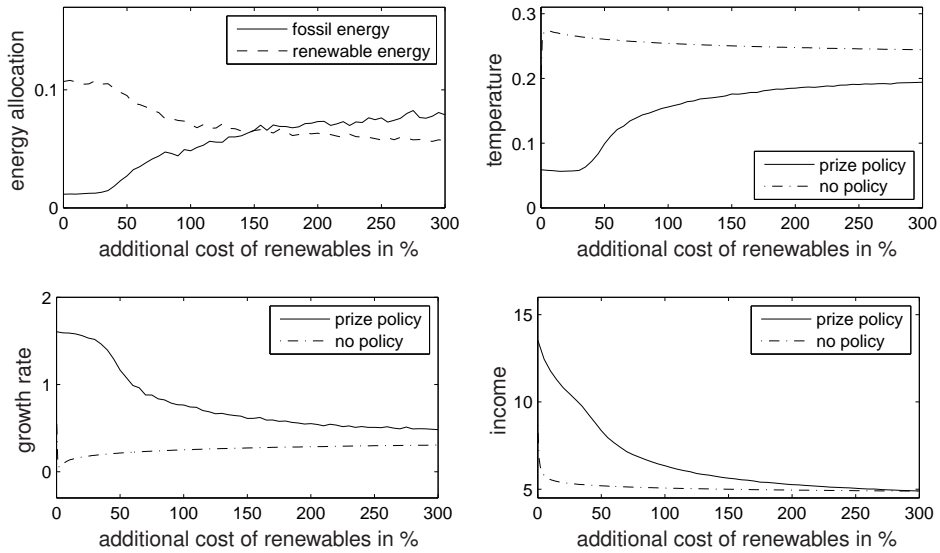

FIGURE 11. Effect of the *prizes* policy for different levels of cost of renewable energy, at $t = 800$. The number of rewarded agents is $q = 3$, the environmental tax on income is $\varepsilon = 1\%$.
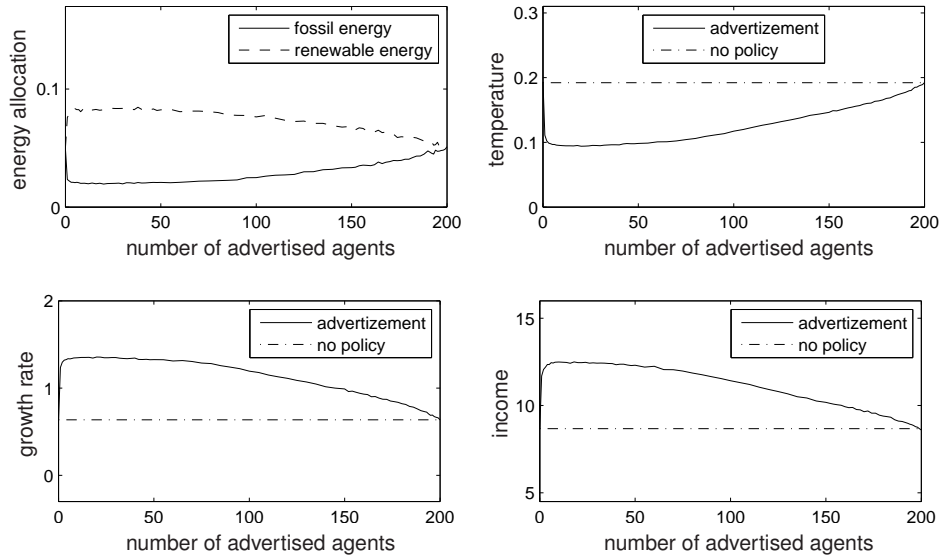
FIGURE 12. Effect of the *advertisement* policy for different numbers of advertised agents, at $t = 800$. The additional cost of renewable energy is $c = 0\%$.
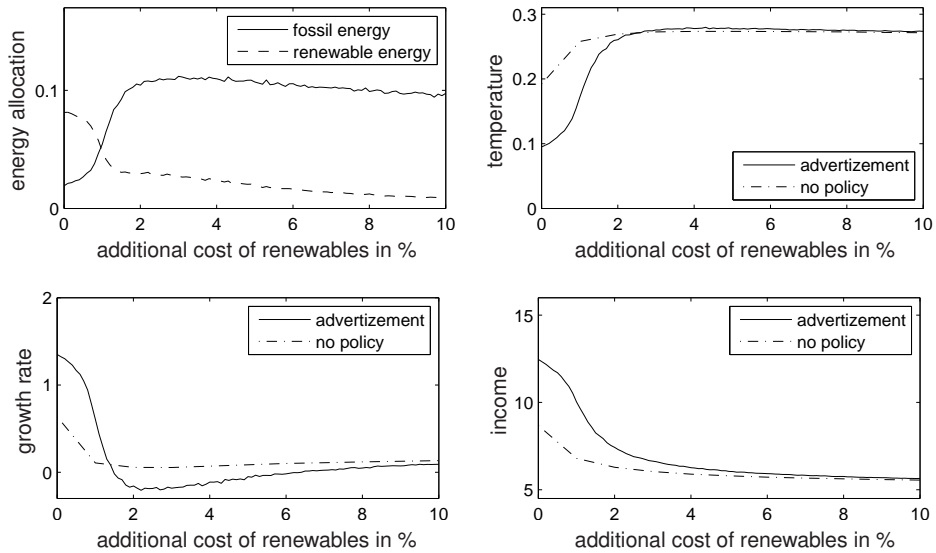

FIGURE 13. Effect of the *advertisement* policy for different levels of cost of renewable energy, at $t = 800$. The number of advertised agents is $q = 10$.