Agent-Based Computer Economics as a Tool for Dialogue between Christian Ethics and Microeconomics

Second Draft – April 2009

Raúl González-Fabre, S.I.
Universidad Pontificia Comillas
Madrid (Spain)
rgfabre@upcomillas.es

Abstract

Starting from a diagnosis of the limitations of rational choice theory to model human ethical behavior [1], the paper proposes Agent-Based Computer Economics (ACE) as an alternative in order to reintegrate ethics within microeconomic analysis, and thus to restore the dialogue between Christian ethics and microeconomics in the very core of the later [2]. For that purpose, ACE methodology is presented and its main uses and limitations are discussed [3]. A strategy is proposed for building ACE models with a theoretical aim, where the proportion of agents who decide according to ethical principles different from utility maximization is the main parameter [4]. Finally, an example is developed where the same microeconomic situation is modeled with only neoclassical agents and with variable proportions of agents that make their decisions according to a principle of transactional justice proposed by the Scholastics [5].

1. The Question

Neoclassical microeconomics rests on a number of assumptions regarding the behavior of the economic agents. Some of them are related to their cognitive capabilities, some others to the way they structure their preferences and make decisions. Given that those decisions are supposed to be intentional, free and intelligent, they belong, without doubt, to the moral life of the agents. Neoclassical economics actually implies an ethical vision of the agents in the economy, eventually disguised as a description of human 'rationality'.

That ethical vision is expressed in an order of preference over all possible states of the world that can be achieved by the agent's decision. In most neoclassical models, those states of the world differ from each other only in the basket of economic goods, services and assets in possession of the decision making agent. In some other models, the states of the world under consideration may also include goods, services and assets in possession of other agents (the economics of altruism and envy). In any case, however, the approach is consequentialist, and the consequences considered are limited to the possession of economic goods, services and assets.

The agent's order of preference, sometimes represented by a utility function, must be complete and transitive over all possible baskets of the relevant economic goods, services and assets. Although it is not usually made explicit, the order of preference of the neoclassical agents is also supposed to be non-lexicographic. That means that any loss of an amount of one of the goods involved can be compensated by a certain gain in other goods, so that there is no 'sacred' good, no good that is always preferred to any amount of another good.

This ethical vision is sometimes condensed into the figure of the 'Economic Man' (*homo oeconomicus*), who may not be egoistic (though he generally is) but it is always a non-lexicographic consequentialist over a limited range of consequences of his decisions. Surprisingly, to say the least, this kind of agent is labeled 'rational'. On the other hand, when neoclassical economics comes to evaluate a certain economic arrangement, such as a market structure, or a policy proposal, its main variable for evaluation is efficiency. Economic efficiency is calculated as the net sum of the utilities achieved by all the agents involved, if those utilities are considered somehow comparable, or it is defined in terms of Pareto optimality in other case.

In the ethics underlying neoclassical economics, moral principles are 'externalized', so to speak, to the civil law. Neoclassical models, except when dealing with the economics of crime and the law itself, presuppose that all economic agents abide equally by the law. They all respect private property, fulfill contracts, pay taxes and obey regulations. The rules are the same for all agents, and all agents act in the same way with regard to them. The law thus can be understood as a system of external constraints on the agents. Those constraints may also be evaluated in terms of the efficiency of the economic dynamics they provoke or allow.

We have, then, that from a moral point of view neoclassical economics composes (1) agents that are homogeneously non-lexicographic maximizers of economic consequences with (2) an external set of compulsory rules that embodies all moral principles deemed relevant to the situation. As a consequence, moral variability in the agents is put aside, and individual ethics cannot be easily incorporated into neoclassical models. On the other hand, given that efficiency understood as the maximization of utility is the main evaluative category both for individual behavior and for the system's dynamics, it takes precedence over other possible normative categories, most remarkably over the fairness of the transactions and the social justice of the whole arrangement. It is perfectly possible to author a book of intermediate microeconomics without ever using any of those two words.

Neoclassical microeconomics has been developed according to the positivist ideal of a 'social physics', even if far less successfully than expected with regard to its predictive accuracy. That positivist ideal requires that models may be mathematically formalized and solved with the tools of algebra, calculus, topology and the like. Supposing that all agents are maximizers allows walking quite a long way in theoretical development using calculus. In our opinion, this fact is what has fidelized standard microeconomics to its moral assumptions, rather than a belief in the realism of those assumptions, or any empirical validation of the results obtained. That theoretical fidelization is remarkable, for it has been kept even when other assumptions of neoclassical microeconomics, such as the absence of 'friction' in the transactions or the 'perfect knowledge' of the agents, were relaxed. Most of current economics, neoclassical or not, share the same fidelity to the moral view of economic life contained in neoclassical microeconomics.

Obviously, this has important cultural consequences. Basic microeconomics, with its underlying moral assumptions about the person in society, is taught in secondary school and forms part of the curricula of many college and university programs. The same assumptions are present not only in the neoclassical economics learnt and used by economists, but also in many other more or less heterodox approaches. And they are quickly expanding into other social sciences, as neoclassical, game theoretical and similar models are used to formalize family, tribal, organizational and political relationships.

As a consequence, the decision making algorithm of the Economic Man is increasingly being sanctioned as the standard of human behavior, the way most people act and the way one should act in economic life or, more generally, in social life. Acting otherwise would be non-rational, would lead to individual or social underperforming, we are told. Thus, that algorithm is becoming a true moral paradigm, used to understand interpersonal and impersonal relationships, and to understand ourselves as agents. Given that the way we interpret others' behaviors and understand our own rationality has a strong influence over the direction in which we consider morally adequate to proceed, the moral assumptions of neoclassical microeconomics tend to be self-fulfilling.

This poses a major challenge to Christian ethics. Here we will include under 'Christian ethics' any ethical body of thought compatible with the teachings of the Gospel. There are many such possible bodies of thought, but all of them share several characteristics that make them incompatible with the assumptions of neoclassical microeconomics. On the one hand, in Christian ethics there is always some kind of moral law independent from and superior to the provisions of civil law. That moral law has a lexicographic structure: some valuable entities (for example, human life) or relational qualities (for example, justice) have strict precedence over others (for example, wealth or efficiency) in all situations of choice where the alternatives differ from each other in the advancement of those values.

On the other hand, in Christian ethics agents are known to be heterogeneous in their fulfillment of the moral law. They possess different moral qualities. Apart from this moral variability across the population of agents, there is also a possibility of moral change over time in the agents. They can improve --or worsen-- their moral quality as a result of both progressive (habit building) and sudden (conversion) internal changes.

Finally, Christian ethics consider relevant not only the external consequences of human actions, but also their internal consequences, such as the changes in the moral qualities of the agent, her relationships and social life in general, produced by those actions. Different approaches in Christian ethics give a different weight to the expected external consequences of the action in the decision making process. It is possible to propose a 'categorical imperative' approach to ethics within the Christian tradition, where moral principles are all and external consequences are of little relevance. It is also possible to follow a more 'prudential' approach, including external consequences in the Christian discernment along with moral principles and internal consequences. In the Christian Western tradition, the former position is more often found in thinkers who use Platonic or Kantian anthropological bases for their ethical doctrines, while the latter is more often found in thinkers who draw their moral anthropology from Aristotle or St. Thomas Aquinas.

The neoclassical moral profile of the economic agents sketched above is obviously incompatible with the vision of Christian ethics in all relevant points. As a consequence, Christian ethics has had historically very little to say in the development of modern microeconomics. Its only possible theoretical role within the neoclassical framework would be to propose sets of compulsory rules external to the markets. Those rules would impose on the agents' behavior some moral principles (fairness) and/or would promote some policy objectives at the aggregated level (social justice). This is theoretically most unsatisfactory, given that economic action is definitely moral action and must be understood as such, not as the basic unit of a quasi-physical dynamics that can only be governed from outside.

In fact, within the Scholastic tradition we can find an important precedent of integration of the Christian ethical vision into the decision making process of the economic agents. Following Aquinas, from XIIIth to XVIIth centuries Scholastic theologians and moralists analyzed all relevant contracts of their time from the point of view of fairness, proposing a consistent lexicographic priority of fairness (*iustitia*) over profit in the agents' decision making. Apart from contributing relevant insights about the dynamics of a market economy that were used by later political economists they produced nuanced theories about the adequate composition between economic consequences and the requisites of fairness in particular market situations.

Unluckily, the creativity of the Scholastic tradition had died down at the end of the XVIIth century. Nobody was there to propose the incorporation of moral issues into the decision making process of the agent in the political economy born during the XVIIIth century. Nor was anyone there to do the same job when, at the end of the XIXth century, the neoclassical revolution took place. Such concerns as Marshall's[1] didn't find a way of incorporation into the core of the new microeconomics.

The ultimate question this paper deals with is, thus, whether the intellectual project of the Scholastic tradition in economics can be resumed, so that the vision of the economic agent proposed by Christian ethics may find a place in the foundations of microeconomics. From a theoretical point of view, that would contribute some moral realism to microeconomics, hopefully increasing the accuracy and relevance of its predictions. It would also restore the dialogue between Christian ethics and microeconomics, necessary given that both deal with intentional, intelligent and free actions of people. From a practical point of view, it would allow Christian social thought to estimate the economic consequences of its ethical proposals, without having to deny itself through the choice of the instrument for such estimation. In the current state of the science, when Christian thinkers propose a certain course of action to some agents, they must either disregard the economic consequences of taking that course, calling to an action based solely on moral principles, or estimate those consequences using a microeconomics based more or less closely on the neoclassical ethical paradigm.

Our ultimate question can be split into two: (1) Finding decision making algorithms for the economic agents compatible with Christian ethics; (2) Finding a methodology for modeling microeconomic situations with variable proportions of agents that make their decisions according to those algorithms, and agents that follow the neoclassical model of decision making. This methodology should allow estimating the economic consequences of the presence of 'Christian-ethics' agents for themselves, for the rest of the agents and for the aggregated system.

This paper deals with the second question. It suggests that Agent-Based Computer Economics is an adequate modeling tool for the purpose of reintegrating Christian ethics into the core of microeconomics.

2. A Possible Answer

The point in question is whether a microeconomic situation with agents of two classes, similar in all regards except for their moral profile, can be modeled and the model solved. The first class of agents of our interest is neoclassical ones, non-lexicographic maximizers of economic results for themselves within the constraint of the (civil) law. The second class is 'Christian' agents, who maximize economic results for themselves conditionally to the fulfillment of some moral rules of fairness, to which nobody is legally constrained. For this second class of agents, decision making follows a lexicographic preference order, where fairness always is preferred to profit.

If such models can be built and solved, nothing would prevent us from evaluating their results giving lexicographic precedence to social justice over efficiency. This would be in accordance not only to the Christian program (justice is due, efficiency is not) but also, for instance, to Rawls' (1999: 3):

[1] "But ethical forces are among those of which the economist has to take account. Attempts have indeed been made to construct an abstract science with regard to the actions of an "economic man," who is under no ethical influences and who pursues pecuniary gain warily and energetically, but mechanically and selfishly. But they have not been successful, nor even thoroughly carried out. For they have never really treated the economic man as perfectly selfish." http://www.econlib.org/library/Marshall/marP0.html. Maybe Marshall would have a different opinion about the current state of standard microeconomics.

> Justice is the first virtue of social institutions, as truth is of systems of thought. A theory however elegant and economical must be rejected or revised if it is untrue; likewise laws and institutions no matter how efficient and well-arranged must be reformed or abolished if they are unjust.

Our answer to this methodological question is that Agent-Based Computer Economics (ACE) offers a suitable tool for the purpose of building microeconomic models with morally heterogeneous agents. The rest of this epigraph is devoted to the basics of ACE models. In the next epigraph we will discuss some limitations of the ACE methodology. We will conclude with a rudimentary exercise on how ACE could be used to estimate the economic consequences of the presence of 'Christian-ethics' agents in a certain, simple economic setup.

Agent-Based Computer Economics is an application of Agent-Based Modeling to economics. Its development has been parallel to the growth in the capacity of computers to run more and more complex models. ACE models are dynamic representations of an economic situation. They consist of two basic components: the environment and the agents. The environment is made of all non purposeful elements of the model—a certain space, flows of matter and energy, external rules that apply mechanically...—while the agents are its purposeful elements. They interact with other agents and with the environment trying to achieve certain objectives that can be represented with a utility function.

ACE models are usually built as object-oriented computer programs. Both the elements of the environment and the agents are instances of certain classes characterized by some state variables and some procedures or methods that change those state variables using information both internal and external to the object. The agents' procedures include their strategies to try and achieve their objectives. The objects that represent agents are often endowed with some learning capabilities that allow them to change not only their state variables but also relevant aspects of their strategies, according to their reading of the local situation and to the individual results they obtained previously.

Time in ACE models is generally discrete, not continuous. The programmer establishes an initial setup and designs an order of the operations that the objects in the system are going to undertake. Those operations are repeated in a loop for a certain number of cycles (ticks). Each tick represents a time interval (one hour, one day, one year...) where the objects develop the operations they were programmed for. The situation at the beginning of the first tick is the initial setup established by the programmer. The beginning situation of the following ticks is the ending situation of the previous tick.

ACE is different from other methodologies of dynamic economic simulation in that no equations are posed and solved for the whole system. On the contrary, the values of the aggregated variables at a given point in simulated time result from the local, decentralized interactions between the individual elements of the model. ACE models build their dynamics 'from the bottom up'.
This makes ACE particularly suitable to model bounded rationality situations, and in general any dynamic economic situation where the corresponding system of differential equations would be very difficult to pose, or mathematically intractable once posed.

Each agent in an ACE model is a distinctive computational object with its own characteristics (variables and procedures) that must be designed in full detail and that can evolve in time without making the model incur calculus intractability, given that no general equations are to be solved. The suitability of ACE for modeling moral heterogeneity in a population of economic agents derives from this. If the model contains all the variables relevant to the agent decision making, there is no calculation limit to the complexity and nuances of the algorithms the agents may use to make their decisions. It is not more difficult to program a lexicographic order of preference over alternatives than a non-lexicographic one. And it isn't more difficult either to program a model with morally heterogeneous agents than to do it with morally homogeneous ones. It suffices to add as many subclasses of agents as moral profiles are to be incorporated to the model.

This way, moral heterogeneity across populations, moral variability over time, and complex lexicographic algorithms for decision making can be easily introduced into microeconomic models. This is done at the level of the very foundations of those models: in the agents themselves. The basic understanding of Christian moral anthropology becomes no longer incompatible with microeconomic analysis, because the latter's assumptions that originated the incompatibility can be removed in ACE models. As a consequence, connections of special interest for Christian ethics, such as the one between the fairness in economic transactions of the agents and the justice of the social arrangement can be analytically studied using microeconomics.

3. Theoretical Scope and Limitations of the ACE Methodology

ACE models can be placed midway between theoretical analysis and empirical research. Designing the environment, the agents and their interactions in time is a theoretical task. As in all theory, the always-too-complex reality must be simplified in order to separate some aspects for study. Options have to be made regarding relevant variables and adequate algorithms.

On the other hand, the results of the simulation can be dealt with using the same tools as empirical research, including econometric tools. Usually ACE simulations are run several times within a certain parametric space, varying the values of the relevant parameters in a discrete fashion. Each run can be understood as a computer experiment. The set of results can be processed statistically in order to extract some conclusions.

But, what is the value of those conclusions? Standard microeconomic analysis based on calculus and topology is able to achieve certainty in its conclusions at the price of great simplification in the design of its models. A small parametric space, with few independent variables, is exhaustively explored. The conclusions are presented in the form of theorems than can be proven.

At the opposite extreme, empirical observation offers a fully detailed particular case for inductive reasoning. Certainty can never be achieved by means of induction from particular cases, and empirical experimentation in real, complex contexts is usually difficult in social sciences. Very often a few cases are all we have as a starting point for induction in social sciences.

ACE stays halfway between microeconomic analysis and empirical observation in this regard. On the one hand, ACE models make a number of simplifying assumptions but they need not be as strict as in analytical models. One can explore the variation of more parameters and aspects of the model, given that calculus tractability is not a concern.

On the other hand, ACE models produce 'experimental' results on the basis of those assumptions. Unlike empirical observation, the number of experiments that can be done is only limited by the available time and computer power. The results of an ACE model are not expressed by means of theorems or mathematical formulae, as in microeconomic analysis, but in a statistical fashion similar to that of empirical research, with the only difference that more experiments can be easily added if necessary to clarify a certain point in discussion.

ACE models may be used for at least two different purposes. The first one has a theoretical intention. More than predicting, these models intend to explore and/or explain social dynamics, showing the underlying mechanisms that (may) generate certain systemic results from local interactions. Research along this line has focused on emergent macrosocial phenomena, often interpreted using categories taken from the theory of chaos and related theoretical corpuses.

The second possible purpose is rather practical. A real situation is modeled in order to find out some likely evolution scenarios from a certain moment onwards. All observable parameters of the model are quantified using empirical data (deterministic or stochastic). The model produces predictions about the evolution of the situation when different events happen, or different courses of action are taken. Those predictions can be used for decision making in policy and management.

Models with an empirical basis and a predictive intention can be fine-tuned (or, using neural networks language, trained) using past empirical data, and they can be validated by comparing the results of the model with the actual evolution of the modeled situation. That way, a reality check is done on the model.

Validation is much more problematic when dealing with models with a theoretical purpose. In this case, the value or interval of values for each parameter of the model must be set by the programmer, according to her theoretical aim, that is, to the social space she intends to explore. Running the model for a continuous interval of values of any one parameter is not possible, as it would imply an infinite number of runs. A choice must be made regarding which finite set of values of each parameter are going to be explored. The 'curse of dimensionality' forces the programmer to be restrictive: exploring ten possible values of each of ten independent parameters would require 1E10 runs of the model, as that is the number of different sets of parameters.

Moreover, most ACE models are built using procedural programming. Both the content of the instructions and the order of their execution are deterministically fixed by the programmer, so that at any given instant, one and only one instruction is being executed. Both the discrete passing of time and the procedural nature of the program limit the realism of ACE models, for in real life economic agents act in a continuous, parallel fashion. Pseudo-random numbers are used to minimize the impact of these limitations. When the order in which the agents or the objects of the environment must act is relevant to the overall result of the model, that order is randomized. That way, the advantage or disadvantage that a certain agent may gain from being first at a certain tick is compensated by the fact that the same agent is in a different position in the order of execution in previous and future ticks. But this solution to the question of procedurality poses an additional problem: for each set of parameters, several runs have to be made in order to ensure that the random order in which the agents are executing their instructions does not create a computational 'artifact'. If 10 runs are made for each of the 1E10 sets of parameters of the above paragraph, we would already be committed to 1E11 runs of the model.

Other computational artifacts may arise from the order of the operations established by the programmer, from the inaccuracies of floating-point computer arithmetics, and from plain programming mistakes. In all cases, as theory-oriented models do not produce predictions that may be checked against real situations, it is difficult to see how those artifacts may be detected. There is a risk that they may pass as genuine results of the model run.

This risk is even bigger given that ACE simulations can be properly revised only by reading the source code. Even if ACE programmers were to make available their source code along with the results of their research, something that does not always happen, that code could be written in any of many different programming languages and simulation environments. There is no standard programming language or simulation shell to write ACE models. This makes it difficult for the average reader to check the model for mistakes and computational artifacts.

However, if the model is well described, it should be replicable in any other language or shell, arriving at the same results. On the other hand, models written in a modular fashion (using object-oriented programming or otherwise) allow the programmer to reuse well-tested chunks of code. This reduces greatly the risk of programming errors.

In summary, the reliability of ACE models with a theoretical aim depends greatly on the carefulness of the programmer and the replicability on a different programming platform of the results she publishes. This forces her to provide either the source code or a very detailed specification of the flow of the program.

Once the reliability and replicability of the results are ensured, what we obtain from an ACE model is a number of computer experiments in an n-dimensional space, being n the number of independent parameters of the model. Each experiment describes the dynamics of the system for a certain set of values of the parameters. The results take the form of a time series for each variable deemed relevant by the programmer. Statistical induction can be attempted on those time series, in order to reach conclusions of a theoretical nature about general trends and counterexamples of those trends, states of equilibrium and other attractors, sensitivity of the system to changes in each parameter, phase transitions and emergent phenomena in general. As indicated above, the easiness of adding more experiments facilitates the exploration, making it possible to ensure the robustness of the conclusions.

4. Using ACE Models with Christian Ethics: a Strategy

Both the theoretical and the practical approaches of the use of ACE models we mentioned above are interesting for our objective of re-inserting Christian ethics into the foundations of microeconomics. The relevance of the practical approach is obvious: if we manage to model adequately a real situation, we can offer moral advice to the involved agents with some approximate knowledge of the expected economic consequences of following it. Including that estimation of consequences, our discernment will gain in prudence.

This paper is more concerned with the other approach, the theoretical one. In order to reinsert Christian ethics into the foundations of microeconomic theory, we propose the following strategy to start with: Using a good ACE model of an economic situation where only morally neoclassical agents are present, we can introduce a variable proportion of agents that follow a 'Christian-ethics' decision making algorithm. Comparing the results of the two models, without and with 'Christian-ethics' agents, we may reach some conclusions about the direction in which the presence of the latter influence the outcome of the model. We may even find emergent phenomena provoked by that presence, and relate all of it to the proportion between neoclassical and 'Christian-ethics' agents.

So, the main intention of this use of ACE models would be a comparative one: to answer the question about the impact of the presence in an economy of a variable proportion of 'Christian-ethics' agents, along with morally neoclassical ones. For the purpose of accumulating knowledge in that direction, models have to be chosen where 'Christian-ethics' and 'neoclassical-ethics' agents would act differently in the situation given. In those models, the most relevant parameter would be the proportion between 'Christian-ethics' and 'neoclassical-ethics' agents. When that proportion is zero, we have a fully 'neoclassical-ethics' model; when the proportion is infinite, we have modeled a utopian situation. In between, we'll be able to find the cases of our interest.

An advantage of this comparative-theoretical approach is that we can start from ACE models of situations where an analytical solution is possible when all agents are of the neoclassical type. That way, we would have a set of theoretical results with which to validate our first model, overcoming the difficulty presented in the previous epigraph. Once validated, that code could be reused to model more complex situations intractable with standard analysis. Further modifications of the model to incorporate 'Christian-ethics' agents would only require extreme carefulness in a few additional chunks of code.

This strategy is to be developed in three steps:
1. Design a simple microeconomic situation where the decision making process of 'neoclassical-ethics' agents would differ from that of 'Christian-ethics' agents. Solve that situation for only neoclassical agents using standard microeconomic analysis.
2. Represent the same situation, with only neoclassical agents, in an ACE model. This implies to express the dynamics of the model in terms of local interactions between agents with bounded rationality. Validate the results, and this the design and codification, of the ACE model comparing with the expected results analytically obtained in step 1.
3. Carefully introduce the necessary changes in the ACE model to allow for a variable proportion of 'Christian-ethics' agents. Run the new model for different proportions of 'Christian-ethics' agents and compare the results with the neoclassical-only version, being all other parameters the same. Study the impact of the presence of 'Christian-ethics' agents on themselves, on other agents, and on the overall dynamics of the model.

5. An Example

In order to demonstrate the above strategy, we propose here a simple, tractable microeconomic model derived with modifications from Sprigg & Ehlen (2004). The model represents a two-market closed economy composed of F firms and H households. There is a labor market and a goods market. Firms buy labor from the households, use it to produce with a fixed productivity, and sell the goods produced to the households. Households have different capacities to produce the same goods at home, so that they must choose whether they will do so or rather seek employment for a wage in a firm.

Let us go through the three steps mentioned above in a summary fashion:

5.1. Standard microeconomic analysis

5.1.1. Households

The households attempt to maximize their utility u_h:

$$u_h = q_h \tag{1.1}$$

where q_h is the amount of goods available to them at each cycle.
The goods are perishable. The households consume the available amount.
The households have two choices in order to acquire goods. They may produce them at home or sell their work in the labor market and buy goods in the goods market with the money earned. The productivity of the households is $\rho_h : \rho_{h\min} < \rho_h < \rho_{h\max}$, variable from one household to another within the indicated limits.
Each household is endowed with a labor capacity $l_h = 1$. If she produces at home, the amount produced is:

$$q_h = \rho_h l_h \tag{1.2}$$

If she rather sells his work to a firm, she will a get a wage w_h. If she can obtain goods at a price p_h in the goods market, she will be able to get for her work:

$$q_h = \frac{w_h}{p_h} \tag{1.3}$$

In order to maximize her utility, each household will concur to the labor market if and only if:

$$\frac{w_h}{p_h} > \rho_h \tag{1.4}$$

In other case, the household will find more profitable to produce at home.

If the household is willing to work for a wage but she finds no employer, we'll say that she is underemployed. Then, she will produce at home due to the lack of a better option.

5.1.2. Firms

The firms attempt to maximize their profit, which is the difference between the income they obtain for the sale of their products and the costs of hiring workers:

$$\pi_f = p_f q_f - w_f l_f \tag{2.1}$$

The firms produce according to this function:

$$q_f = \rho_f l_f \tag{2.2}$$

where ρ_f is the productivity of the firm, which takes the same value for all of them, and l_f is the number of worker hired by the firm at the wage w_f

Putting together (2.1) and (2.2), we get the following condition for the profit to be positive:

$$\pi_f > 0 \rightarrow \frac{w_f}{p_f} > \rho_f \tag{2.3}$$

5.1.3. Market results in monopoly and perfect competition

Given that the households buy goods using only the money they earned in the labor market, the relevant price for this model is the relationship *w/p*.

In the case of monopoly, F=1, it can be analytically proven that the maximum profit is reached when:

$$\frac{w}{p} = \frac{\rho_f + \rho_{h\min}}{2} \tag{3.1}$$

For that price, the number of employed households is:

$$L = H\frac{\rho_f - \rho_{h\min}}{2(\rho_{h\max} - \rho_{h\min})} \tag{3.2}$$

The number of underemployed families is:

$$U = H\frac{\rho_f - \rho_{h\min}}{2(\rho_{h\max} - \rho_{h\min})} \tag{3.3}$$

and the total production of the firms per unit of time:

$$Q_F = \rho_f L = \rho_f H\frac{\rho_f - \rho_{h\min}}{2(\rho_{h\max} - \rho_{h\min})} \tag{3.4}$$

In the case of perfect competition, $F = \infty$, the corresponding analytical results are:

$$\frac{w}{p} = \rho_f \tag{3.5}$$

$$L = H\left(\frac{\rho_f - \rho_{h\min}}{\rho_{h\max} - \rho_{h\min}}\right) \quad (3.6)$$

$$U = 0 \quad (3.7)$$

$$Q_F = \rho_f L = \rho_f H\left(\frac{\rho_f - \rho_{h\min}}{\rho_{h\max} - \rho_{h\min}}\right) \quad (3.8)$$

5.2. ACE Simulation

5.2.1. Parameters and expected results

The second step in our methodology is to design and run an ACE model of the same situation we have just solved with standard microeconomic analysis. All the parameters in the model have to be assigned values. Here we shall explore the following values:

Table 1

H	200
F	{1, 5, 9}
$\rho_{h\min}$	50
$\rho_{h\max}$	250
ρ_f	200

Then, we shall be running three different simulations, one for each initial value of $F = \{1, 5, 9\}$. Each simulation will run for 2000 ticks (cycles), each tick representing one working day. The simulation has been programmed in Matlab. The code will be run procedurally, one instruction at a time. For that reason, a pseudorandom generator will be used to establish the order in which the agents execute their instructions. To prevent artifacts created by that random order, each simulation will be run 10 times with different seeds for the randomizer. The result of the simulation will be the average of those 10 runs.

The expected results, according to the microeconomic analysis of the previous epigraph, are:

Table 2

	monopoly	perfect competition
Relative wage *w/p*	125	200
Employment L	75	150
Underemployment U	75	0
Total production of the firms Q_F	15,000	30,000

5.2.2. Model setup

The basic model consists of 200 households and *F* firms, to be created in the setup stage of the program. The physical distribution of households and firms is not relevant to their interactions. For that reason, there is no positional data associated to them.

The relevant parameters for the households are:

- Their productivity when producing at home: ρ_h. This value is randomly assigned to each household at the moment of her creation, from a uniform distribution in the interval (50, 250).
- The number of firms they are going to visit when seeking a job and when looking for goods to buy. These values are the same for all households: 3 for both parameters.

The relevant parameters for the firms are:

- Their productivity: ρ_f = 200, the same for all firms.
- The initial amount of money available to each firm for hiring workers = 1E5 / *F*.
- The initial price they will ask for the goods: *pIni* = rand * 0.5 + 0.5, the same for all firms. Here rand is a random real number from a uniform distribution in the interval (0, 1).
- The initial wage they will offer: *wIni* = (rand * (ρ_f - $\rho_{h_{\min}}$) + $\rho_{h_{\min}}$) * *pIni* = (rand * (200 - 50) + 50) * *pIni*, the same for all firms.
- The radius of exploration in the plane $w \times p$: *dStep* = 0.01, the same for all firms.
- The maximum radius of mutation in the plane $w \times p$: *dMut* = 0.05, the same for all firms.
- The probability of mutation in the plane $w \times p$ when no other learning choice is open: *pMut* = 0.25.

It must be noticed that all firms have exactly the same initial parameters.

5.2.3. Model operation

The model runs in 2000 ticks where the following operations are executed sequentially:

1. *Households and firms are initialized for the tick*. All tick-related variables are assigned their default values (usually 0 or []). This includes severing all contractual links between households and firms, so that every household is unemployed at the beginning of the tick, and destroying all remaining stock of goods possessed by the firms, given that the goods are supposed to be perishable. The money, however, is carried on to the next tick.
2. *Households seek employment with the firms*. For that purpose, each household visits 3 operative firms (or all firms if only less than 3 are in operation), randomly chosen. For deciding whether to try and sell her work, the household will get the average price at which the visited firms sell their goods and the maximum wage any of them is willing to pay. If the condition (1.4) is satisfied by those values, the household will try to sell her work to the firms in decreasing order of wage offered. The firms always buy the work offered to them, if they have money enough to pay for it.
3. *Households that haven't sold their work produce at home*. The amount produced is calculated with (1.2). This at-home production is destined to self-consumption. It won't reach the market.
4. *Firms produce with the labor they have bought*. The amount produced is calculated with (2.2).
5. *Households with money buy goods from the firms*. Each household in possession of money chooses randomly 3 firms with stock to sell (if there are less than 3 firms in that condition, the household chooses them all). The household visits the chosen firms in increasing order of the prices they offer, buying as much as possible from them until she has spent all her money.
6. *Households consume their whole holdings of goods*, and get utility according to (1.1).
7. *Firms calculate the profit they have obtained*, by comparing their holding of money at the beginning and at the end of the tick.

8. *Firms adjust their strategies.* Each firm's strategy consists of a vector in the plane $w \times p$. Each firm modifies that vector looking for higher profits. For that purpose, a learning process takes place for each firm, with three possible actions in the following sequence:
   a. Exploration: the firm visits one of the neighboring points in the plane $w \times p$, increasing or decreasing *w* and/or *p* by *dStep* (1%). Eight possible directions of exploration are possible depending on whether *w*, *p*¸ or both, are modified positively or negatively. The firm chooses one of those directions randomly. If an increase in the profit takes place in the following tick, the firm goes on modifying *w* and *p* in the same direction. If no increase of profit is registered, the firm chooses randomly a new direction for exploration.
   b. Imitation: if all eight directions have been explored without a profit increase, the firm imitates the current strategy of the firm with the highest profit.
   c. Mutation: with probability *pMut* (25%), if no other firm has obtained a higher profit, the learning firm mutates independently *w* and *p*, increasing or decreasing them between *dStep* (1%) and *dMut* (5%).

   The three processes are mutually exclusive. If the adjustment of strategy takes place through exploration, imitation and mutation don't happen, and so on.

   However, in all three cases there is a limitation to the learning process. The firm cannot plan a strategy that would necessarily produce her significant losses. The condition for positive profit is (2.3): $w/p > 200$. Allowing for 1% of temporary loss, if as a result of the learning process $w/p > 202$, then the firm increases *p* up to $p = w/200$.
9. *Firm survival is verified*. A firm goes out of business if at the end of the tick her money is less than her *w*, meaning that she won't be able to hire one worker in the next tick.

In steps 2 and 5 some interaction between agents happens where the order of execution is relevant. In those steps, the order in which households act is randomized.

5.2.4. Model results

As mentioned above, for each value of $F = \{1, 5, 9\}$ we run the model 20 times with different seeds of the randomizer. The following results are the average for each tick of those 20 runs. They can be compared with the expected results according to Table 2.

Fig. 5.1. Relative price *w/p*

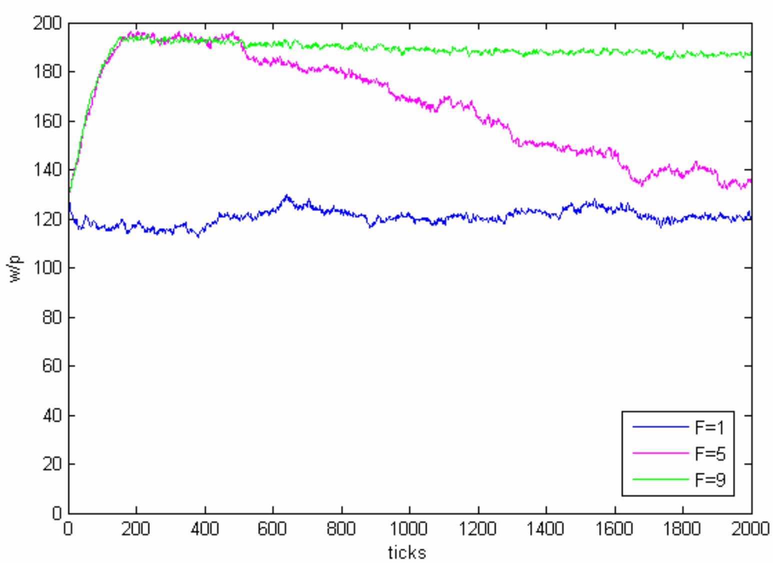

Fig. 5.2. Employment L and underemployment U

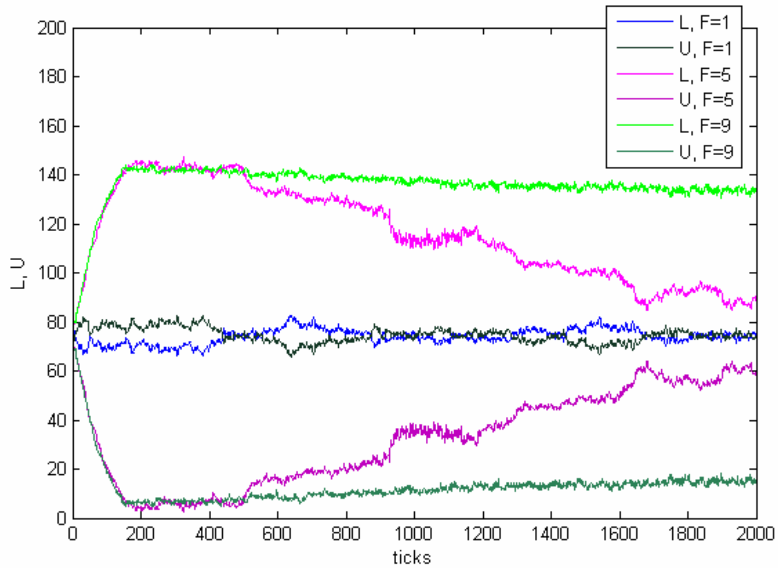

Fig. 5.3. Production of the firms (Q_f) and utility of the households (Σu)

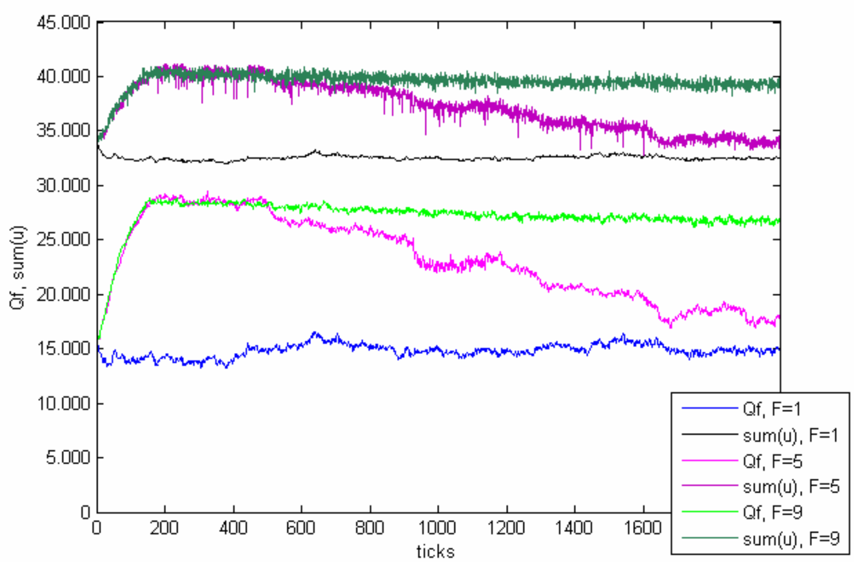

Comparing the graphs with Table 2, we observe a reasonable correspondence between the expected and the obtained results, after tick=190. Before tick=190, a transition develops from the random initial strategy of the firms to the strategies that maximize profits. Both the monopolistic (F = 1) and the perfectly competitive (F = 9) scenarios are well approached. The divergence between the theoretical and the experimental results can be explained by the bounded rationality of the agents. That divergence is greater where the cognitive limitations of the households introduce a bigger gap between the perfect knowledge of the analytical model and the bounded rationality of the ACE model (F = 9). The design and encoding of the ACE model is thus validated.

The intermediate case (F = 5) shows a transition of the results, that start at a perfectly competitive level and end close to the monopolistic scenario. This is due to the exit of business of some firms along the run, which makes the market less competitive.

5.3. ACE Model with 'Christian-ethics' Agents

5.3.1. Model Modification

Once the main code of the ACE model has been validated, we can modify it to transform some firms into 'Christian-ethics' agents. In this most simplified example, we'll limit that modification to the moral prohibition of exploiting monopoly power. That prohibition makes an essential part of the Scholastic theory of fairness (*iustitia*) in the market behavior of economic agents, but it is not a regular part of modern civil law. On the contrary, creating some monopoly power and exploiting it seems to be the

In more complex models, some moral standard of fairness could also be requested from the households, and a reputation mechanism could link fair agents on both sides of each market. Moral variation of the agents over time could also be incorporated.

For our demonstrative purpose, we shall translate the moral prohibition of exploiting monopoly power as follows: the firm that makes an ethical option to abide by that prohibition won't adopt a strategy with $w/p < 180$. That rule has lexicographic priority over the rest of the learning process, aimed at the maximization of profit: whatever the consequences on profit, the rule is to be followed by 'Christian-ethics' agents.

This implies only a small modification in the code of the ACE model: at the end of the learning process of the firm, if this is a 'Christian-ethics' agent, she will verify that her strategy fulfills that rule. If it doesn't, she will adjust it by progressive increases of *w* and decreases of *p*, until it does.

For the rest, the ACE model we have described in the previous epigraph remains unchanged. We then run the model for the following values of the number of firms *F* and the 'Christian-ethics' firms *C:* $F \times C$: {(1,1) (5,1) (5,2) (5,3) (9,1) (9,3) (9,6)}.

5.3.2. Modified Model Results

Let us inspect first the changes that the presence of 'Christian-ethics' agents introduces in the key result of the model: the relative price *w*/*p*. Each graph shows the results for a different value of F. Lines in different colors indicate the evolution of *w/p* for different values of the number of 'Christian-ethics' agents (C):

Fig. 6.1. Relative price *w*/*p*

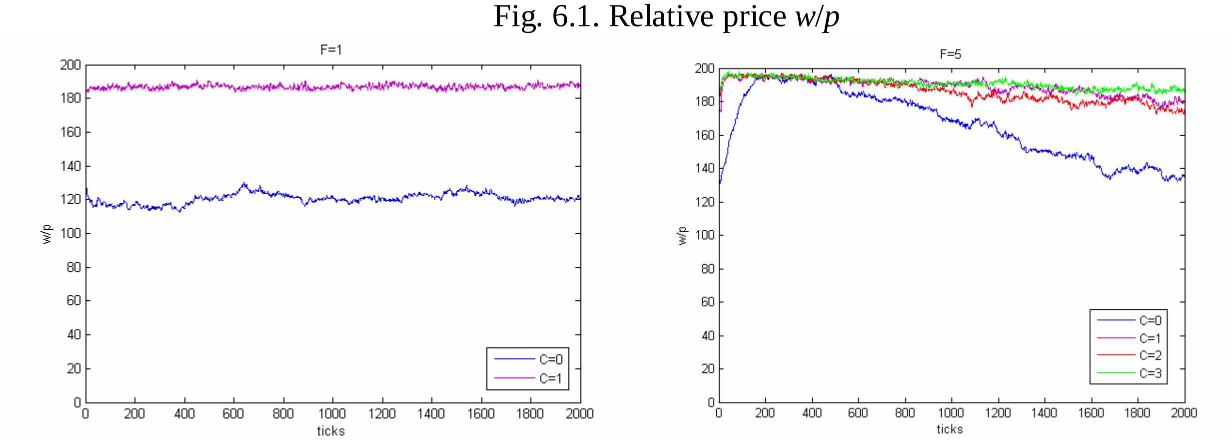

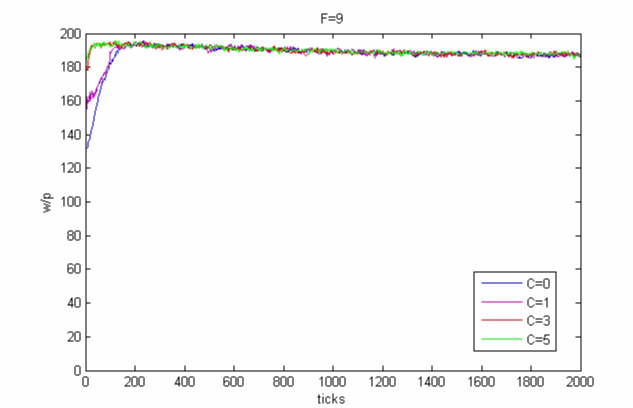

The results are as they could have been expected. In the case of a monopoly (*F*=1), the market depends entirely on the moral profile of the monopolist. If the only firm decides to operate in a 'competitive' rank of prices, the results in term of prices, and in consequence of employment, production and aggregated utility of the consumers, are similar to an actually competitive market.

On the other hand, for a competitive market (*F*=9), the presence of 'Christian-ethics' agents produces very little impact on the relative price. It only makes quicker the transition from the starting point to the competitive equilibrium. The prohibition of exploiting monopoly power operates as a heuristic rule for finding the equilibrium that is going to be reached anyway. This could also be expected: in the case of a competitive market, there is no monopoly power. The prohibition of exploiting it is thus of little relevance.

The most interesting case is that of an imperfectly competitive market (*F*=5). Given that the 'Christian-ethics' firms adjust *w*/*p* using an algorithm different from systematically maximizing profit, it could be expected *prima facie* that they will obtain less profit. With less profit and less money available as a consequence, they would be able to hire less labor than other firms. And with less labor, less production and fewer sales should follow. 'Christian-ethics' agents could thus be in competitive disadvantage, making them candidates for an early exit of the market.

The model results do not confirm these predictions, however. On the contrary, the presence of 'Christian-ethics' firms push *w/p* close to competitive values. If 'Christian-ethics' firms are a majority of all firms (*C/F* = 3/5), *w/p* takes fully competitive values. But even if there is only one 'Christian-ethics' firm (*C/F*=1/5), this firm has a remarkable impact on *w/p*, greater than the marginal impact of any additional firm that adopts the same ethical rule.

Given that this is the main result of our demonstrative model, let us test the hypothesis that the 10 'experimental' values of *w/p* we have obtained for each of C={1, 2, 3} belong to a population with a different mean from the values obtained for C=0 (only neoclassical agents). We'll do it for the values obtained every 100 ticks.

In each case, the null hypothesis H_0 is the equality of the mean value of *w/p* for C=0 and each of C={1,2,3}, at a certain tick = {100, 200,..., 1900, 2000}. We perform a right-tailed two-sample *t*-test of the null hypothesis not assuming equal variances. For a significance level α = 0.05, we obtain:

Table 3

	C=0 vs. C=1		C=0 vs. C=2		C=0 vs. C=3	
tick	conclusion	p-value	conclusion	p-value	conclusion	p-value
100		0.05439	reject	0.04855		0.05224
200		0.70080		0.40164		0.60960
300		0.15122		0.18158		0.29142
400		0.32025		0.45708		0.41588
500		0.18751		0.54976		0.76511
600		0.21360		0.17670		0.22206
700		0.10795		0.11050		0.12349
800		0.12433		0.11328		0.13269
900		0.08994		0.10396		0.06403
1000	reject	0.04116		0.08495	reject	0.03299
1100	reject	0.02823		0.14444	reject	0.02611
1200	reject	0.01200		0.05005	reject	0.02363
1300	reject	0.00539	reject	0.01909	reject	0.00552
1400	reject	0.00586	reject	0.01401	reject	0.00972
1500	reject	0.00296	reject	0.01166	reject	0.00329
1600	reject	0.00580	reject	0.01696	reject	0.00657
1700	reject	0.00030	reject	0.00050	reject	0.00045
1800	reject	0.00024	reject	0.00070	reject	0.00031
1900	reject	0.00038	reject	0.00418	reject	0.00347
2000	reject	0.00066	reject	0.00382	reject	0.00158

In consequence, from around tick = 1000 onwards (1300 for C=2) the null hypothesis can be rejected at the α = 0.05 significance level. An inspection of the p-values shows a considerable robustness of the conclusion: 'Christian-ethics' agents tend to move the main result of the model (*w/p*) to higher values than the ones obtained without them. They tend thus to make the dynamics of the model more competetive.

It is implied, and it can be easily verified looking into the numerical output of the model, that the 'Christian-ethics' firms have better survival opportunities in this model than the neoclassical firms.

For comparison purposes, let us see the other relevant results of the model for *F*=5:

Fig. 6.2. Employment L and underemployment U

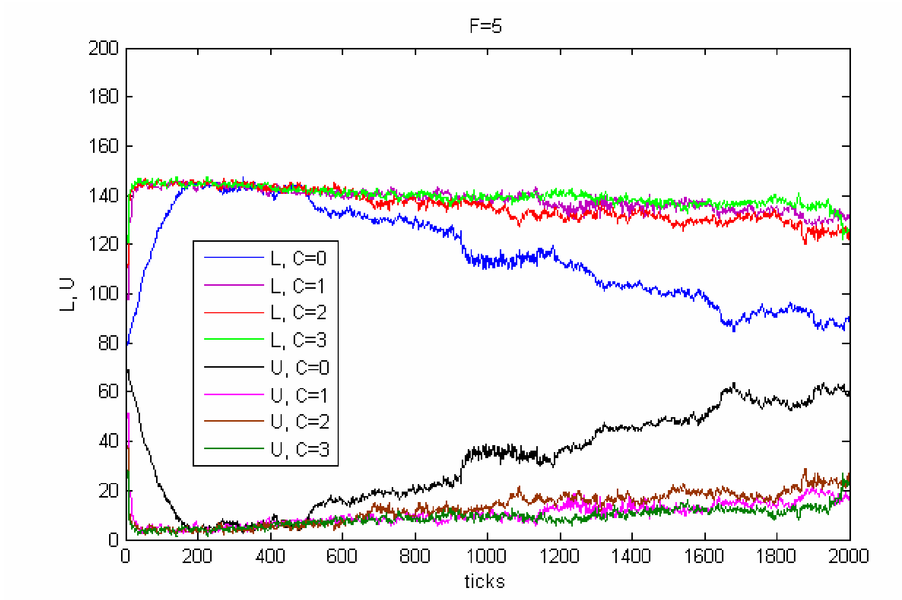

Fig. 6.3. Production of the firms (Q_f) and utility of the households (Σu)

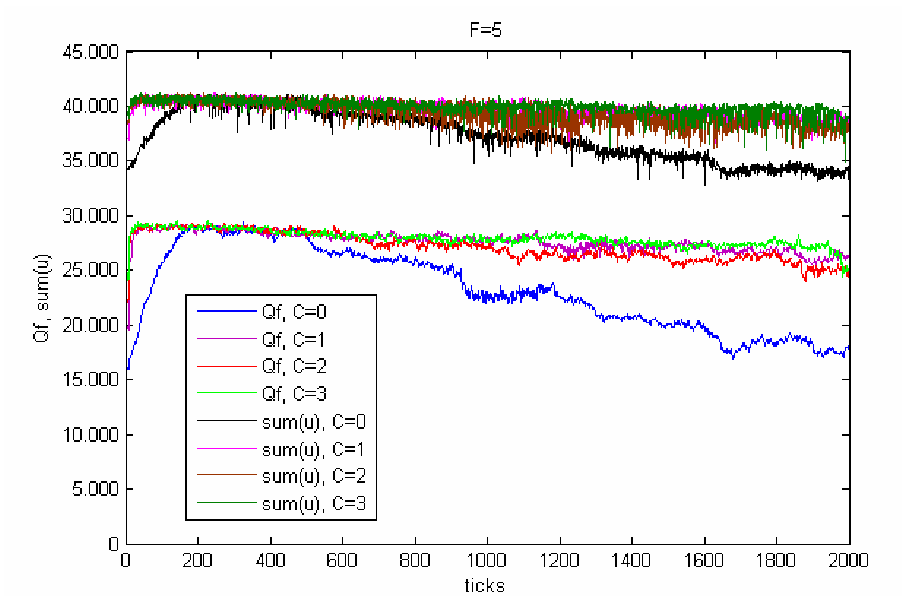

As expected, the results in terms of employment, underemployment, production and aggregated utility follow the evolution of *w/p*. The introduction of even one 'Christian-ethics' agent improves significantly the efficiency and welfare results of the model. When 'Christian-ethics' agents are a majority, the *F*=5 scenario evolves as a fully competitive one.

Finally, we can have a look at the Gini coefficient of the households in each scenario. The Gini coefficient has been calculated on the utility accumulated by each household from the beginning of the simulation:

Fig. 6.4. Gini coefficient of the accumulated utility of the households

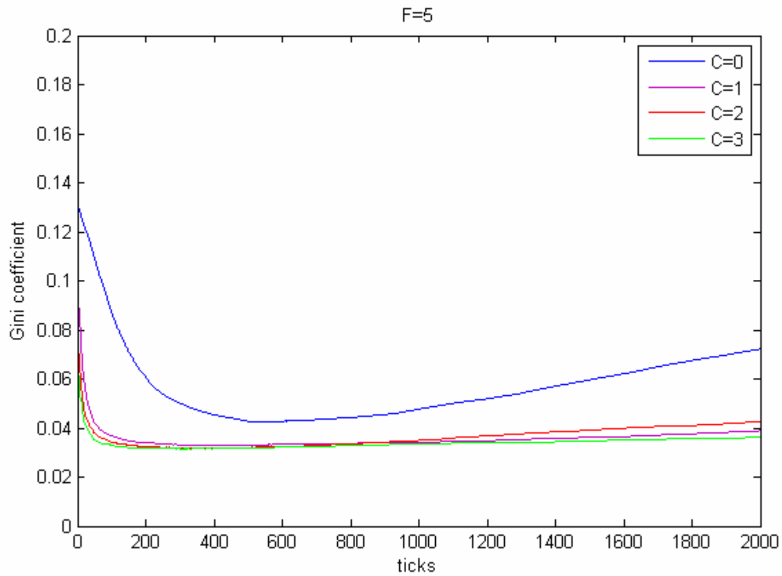

The Fig. 6.4 shows that the presence of 'Christian-ethics' agents in the model tends to reduce the Gini coefficient of the households. It has an equalizing impact.

6. Conclusions

Both the design and the statistical treatment of the model presented above are so elementary, that no actual conclusions can be driven from it. A much more detailed exploration should be carried out in order to evaluate the truth and robustness of the first results. Those results, however, suggest that the old moral rule about not exploiting monopoly power may be sustainable in competitive contexts, even if only one firm follows it. The main effect of the activity of firms following that rule would be to move the market in the direction of greater efficiency and equality.

The model had only a demonstrative intention. Starting from a tractable microeconomic situation, we have been able to reproduce the analytical results in an ACE model. That way, we have validated the ACE code, leaving it ready for use in more complex models. A small, easily controllable modification of that code has allowed us to explore preliminarily the influence of the presence of agents with an algorithm of decision compatible with the Christian vision of moral life, in a model populated also by neoclassical agents.

In summary, we have intended to show a way to reintegrate Christian ethics into the very foundations of microeconomics, using ACE models. From a methodological point of view, two basic challenges remain: to model the essentials of Christian ethics, including the moral variability of people over time, in a way that can be adequately represented in computer code; and to better explore the passage from particular computer experiments to general, robust conclusions.

References and Bibliography

- **Batten, David F.** *Discovering Artificial Economics : How Agents Learn and Economies Evolve*. Boulder: Westview Press, 2000.
- **Bermúdez, José Luis.** *Decision Theory and Rationality*. Oxford ; New York: Oxford University Press, 2009.
- **Blilie, Charles L.** *The Promise and Limits of Computer Modeling*. Singapore ; Hackensack, N.J.: World Scientific, 2007.
- **Browne, Katherine E.; Milgram, B. Lynne and Society for Economic Anthropology (U.S.).** *Economics and Morality : Anthropological Approaches*. Lanham, MD: AltaMira Press, 2009.
- **De Marchi, Scott.** *Computational and Mathematical Modeling in the Social Sciences*. Cambridge ; New York: Cambridge University Press, 2005.
- **Edmonds, Bruce; Hales, David.** "Computational Simulation as Theoretical Experiment," Manchester: Manchester Metropolitan University, 2003, 19.
- **____.** "Replication, Replication and Replication: Some Hard Lessons from Model Alignment." *Journal of Artificial Societies and Social Simulation*, 2003, *6*(4).
- **Epstein, Joshua M. and Axtell, Robert.** *Growing Artificial Societies : Social Science from the Bottom Up*. Washington, D.C.: Brookings Institution Press, 1996.
- **González Fabre, Raúl.** *Ética Y Economía : Una Ética Para Economistas Y Entendidos En Economía*. Bilbao: Desclée de Brouwer, 2005.
- **____.** *Justicia En El Mercado : La Fundamentación De La Ética Del Mercado Según Francisco De Vitoria*. Caracas: Universidad Católica Andrés Bello, 1998.
- **Klos, Tomas.** *Agent-Based Computational Transaction Cost Economics*. Capelle, The Netherlands: Labyrinth Publication, 2000.
- **Marshall, Alfred.** *Principles of Economics*. 1920. Library of Economics and Liberty. Retrieved April 2, 2009 from the World Wide Web: http://www.econlib.org/library/Marshall/marP0.html
- **Mathieu, Philippe; Beaufils, Bruno and Brandouy, Olivier.** *Artificial Economics : Agent-Based Methods in Finance, Game Theory and Their Applications*. Berlin: Springer, 2006.
- **Rawls, John.** *A Theory of Justice*. Cambridge, Mass.: Belknap Press of Harvard Univeristy Press, 1999.
- **Robinson, Stewart.** *Simulation : The Practice of Model Development and Use*. Chichester, West Sussex, England ; Hoboken, NJ: John Wiley & Sons, 2004.
- **Sprigg Jr, James A.; Ehlen, Mark A.** "Full Employment and Competition in the Aspen Economic Model: Implications for Modeling Acts of Terrorism," *Sandia Reports.* Albuquerque, New Mexico, 2004, 28.
- **Tesfatsion, Leigh and Judd, Kenneth L.** *Handbook of Computational Economics*. Amsterdam; Oxford: Elsevier, 2006.