Why Should an Economy Be Competitive?

Hugues Bersini

IRIDIA, Université Libre de Bruxelles
CP 194/6 - 50, av. Franklin Roosevelt
Bruxelles 1050, Belgium

Nicolas van Zeebroeck

ECARES, Solvay Brussels School of Economics and Management
Université Libre de Bruxelles
CP 114 - 50, av. Franklin Roosevelt
Bruxelles 1050, Belgium

A new look at the well-known trade-off between efficiency and equality with an agent-based model is provided. By way of a computer program, the interactions of agents producing and trading goods within different market structures are simulated by our study, which looks at the resulting production and distribution of welfare among agents at the end of an arbitrarily large number of iterations. Two market mechanisms are compared: the competitive market (a double auction market in which agents outbid each other in order to buy and sell products) and the random market (in which products are allocated randomly). Our results first confirm that the superior efficiency of the competitive market comes at a very high price in terms of inequality compared with the random market. The effect of agent rationality in production and auctioning is further explored (i.e., different information sets used or not by the agents in making their choice) and although rationality is observed to affect efficiency only at the margin, inequalities can be very strongly affected by the behavior of the agents. This latter result suggests that market mechanisms can ensure optimal efficiency under certain constraints, but that the degree of inequality emerging from a certain market design can strongly depend on the rationality of the agents.

1. Introduction

Economic tradition holds that competition and free markets are the best way to create value. Adherents to Adam Smith's exposition of the invisible hand and famous historical advocates of the competitive or perfectly concurrent free market, such as Friedrich Hayek, Milton Friedman, and many others, have usually seen this type of decentralized and self-organized economical exchanges between competitive sellers and buyers as the most efficient way to maximize social welfare. More recently, Gode and Sunder [1] have shown that highly competitive market structures such as double auctions lead to efficient product allocation and price setting, even when the agents have zero intelligence (i.e., place offers and bids in a purely random way). They

conclude that imposing a budget constraint on the agents is sufficient to maximize the allocative efficiency of these auctions, even with irrational (or zero-intelligence) agents.

Not surprisingly, a competitive system potentially likely to promote the winners at the expense of the losers may have little chance to equally distribute wealth. Competition in the economy is inherent to the nature of the law of supply and demand that rules the dynamics of these markets, calibrating the product's price such as to adjust demand to supply. As a result of this competition among agents, the price to acquire a rare product should increase, favoring the richest buyers, and the price to sell an abundant product decreases, favoring the highest-performing sellers (who propose lower prices). One classical counterargument to this apparent source of inequality is that, the sellers, by competing, will finally provide products at cheaper prices to the poorest buyers and reciprocally, the buyers, by competing, will allow the unqualified sellers to also profit from the products they sell.

These factors translate into the well-known trade-off between efficiency and equality (see [2, 3]). Although the existence of such a trade-off has been discussed for a while, relatively less attention has been devoted to questions such as: which features of competitive markets trigger or foster inequalities among agents? How does the behavior of agents affect inequality? By giving up competitive market structures, would the loss in efficiency be proportional to the gain in equality? More generally, whereas some scholars have analyzed inequality within competitive market structures, fewer works have investigated the efficiency and egalitarian properties of alternative market mechanisms. Besides, most economists still advocate for matching mechanisms that maximize efficiency and compensate for the resulting inequality through redistributing processes, without questioning much this philosophical choice.

A succession of theoretical models has shown that it is possible to maximize the efficiency in product allocation through price adjustments. However, although all agents benefit from the exchanges at the aggregate level, individual welfare still depends on the agent's tastes distribution, budget constraints, and production skills, and not much is told about the welfare distribution among the agents. Allocative efficiency and Pareto optimality have nothing to do with equality among agents. As a matter of fact, there is a longstanding philosophical dispute between the utilitarians (maximizing cumulated welfare) and the egalitarians (reducing welfare variability) [4], since both objectives (while equally ethically justified) appear to be antagonistic in many social contexts. Clearly, a competitive economy looks toward the utilitarian side.

With regard to the egalitarian perspective, little attention has been paid so far in the literature to the side effects of competitive market structures on welfare distribution among agents. This is a difficult question to address from an analytical point of view. We believe that

conducting such a comparative analysis through a more classical mathematical approach would be very complicated if not impossible with the mathematical instruments usually found in economics. In this paper, a software stylized model is presented in order to compare two styles of market structures (thus purposefully unrealistic): competitive, a double auction market in which buyers and sellers compete by trying to outbid each other, versus random (or distributive), where the matching between buyers and sellers is done in a purely random way insofar as certain constraints are met. The two market structures are compared along two main dimensions: the welfare produced as a result of the agents' interactions (in terms of utility and money) and the degree of inequality in welfare allocation between agents.

The application of agent-based modeling to economic problems is not new, however. Agent-based computational economics, or ACE, is described by Tesfatsion as "the computational study of economies modeled as evolving systems of autonomous interacting agents" [5]. Instead of a top-down analytical approach, this type of model studies economies from the bottom up by simulating a large number of interacting agents and observing what kind of phenomena emerge. Agent-based models are becoming more and more prevalent in economics. A few examples are illustrated in [6–8].

Many mechanisms testify to the clear connections existing between the sort of computational models popular in the complex systems or artificial life literature and the working of free markets. These include the parallel interactions among simple agents (which are just motivated by profit), the reactivity of these agents to the *stigmergic* effect of price, buying less when prices increase and pushing prices to decrease by selling more of the most-wanted products, and the self-organized stabilization of prices that equilibrate supplies and demands. As a matter of fact, John Holland was among the authors of a 15-year-old paper entitled "Artificial Economic Life: A Simple Model for a Stockmarket" [9] where the authors state in the introduction: "This stockmarket model may also be seen as a case-study in artificial life; from a random soup of simple rules an intelligent system spontaneously arises…" As recalled by John Cassidy [10], Friedrich Hayek, the key historical figure of the free market economy, was a sort of complex system pioneer; he stated that market prices are primarily a means of collating and conveying information (just like insects' pheromones) that would be totally impossible to replicate for a centralized cognitive planner.

> Hayek's metaphor of the market as a "system of telecommunications" is more direct and specific. It helps explain how markets work—via the transmission of price signals—and why they are so difficult to replicate. "The most significant fact about this system is the economy of knowledge with which it operates, or how little the individual participants need to know in order to

be able to take the right action," Hayek wrote.... "Only the most essential information is passed on and passed on only to those concerned."

The paper that most closely relates to ours is Gode and Sunder [1], who look at the dispersion of profits among agents in the human versus zero-intelligence markets and conclude that, contrary to the allocative efficiency that can be maximized by the market structure itself, profit dispersion may depend on the behavior of the agents. However, they do not investigate the impact of human behavior on utility dispersion. This paper departs from Gode and Sunder's work in different ways. First, Gode and Sunder look at the difference in allocative efficiency of a single market structure when two different types of agents interact: human agents versus zero-intelligence machine agents who place bids and offers randomly. In contrast, we look at the same sets of machine agents who interact within two different types of market structures: competitive versus random. Second, we not only look at welfare creation (or allocative efficiency), but most importantly at welfare distribution among agents, in terms of both utility and money. Finally, we also investigate the impact of different agents' behaviors in the different markets on welfare creation and distribution.

A computer program will be presented in the following, implementing simple rules for the agents to follow as well as a framework that allows them to interact, while logging facilities collect data on the transactions executed and the evolution of the different metrics in the model. The agents are then set free in numerous simulations, and the resulting metrics are compared across simulations. Our first sets of results show that a competitive market structure such as a double auction consistently leads to much higher welfare at the aggregate level than a random matching market. But they reveal considerable differences in the distribution of utility and money between the two market structures, with the competitive market leading to much more unequal distribution of welfare between agents. We then explore the impact of informed versus zero-intelligence producers. One obvious proof of the poor efficiency of a random production is that an important percentage of simulations just fail since no matching occurs between what the consumers want most and what is being offered by the producers. We find that even in the most competitive market, ignorance of the producers leads to slightly more utility for the consumers but at the expense of much more money being wasted, which therefore makes both markets much less efficient. We finally investigate different behaviors of our machine buyers and sellers in selecting the products for which they want to place bids or offers. We observe that although welfare creation is generally equal, the degree of inequality in utility and money distribution with the competitive market significantly depends on the rationality of the agents.

In Section 2 the classes of the object-oriented model are described. Then the experimental results of the cumulated distribution of wealth

will be compared across the two types of economy: competitive and random. As usual with these kinds of artificial models (such as with classical game theory models and many complex system models), the simulation presented here is not intended to depict any precise reality but needs to be construed as a software thought experiment, namely the conception and the execution of virtual worlds helping to understand in outlines the behavior of a purposefully caricatured reality.

2. The Model

The model implemented in C# maps elegantly to an object-oriented model with the distinct responsibilities distributed through the different classes. These classes and their relationships can be seen in the simplified class diagram in Figure 1.

2.1 The World

The model's different components all live within a structure called the *world*. The world contains all the agents as well as the market. Each world has one market (either a competitive double-auction market or a random market), a series of agents, and some world specific settings, such as the initial endowment of the agents and the number of different products the agents can make and trade. A given number of products are bought and sold. The world is not limited in the number of units for each product, but each transaction concerns only one unit of the product. Each world corresponds to one simulation. Worlds always come in pairs with equal initial settings, where one takes care of the competitive market and the other one takes care of the random market.
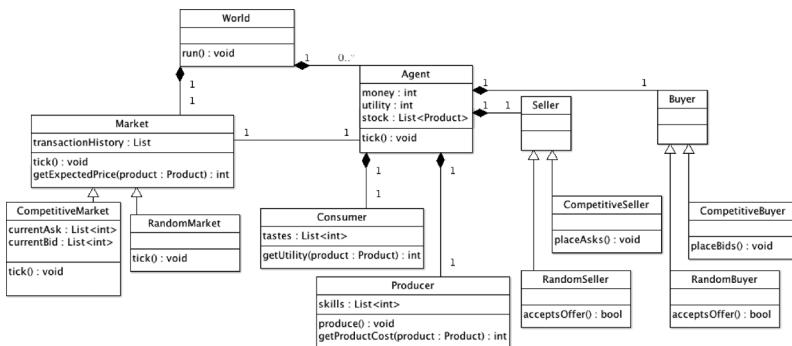

Figure 1. The UML class diagram of the model.

2.2 The Agents

The agents are the main actors of the model; they are the imaginary people who produce, consume, and trade goods driving the model's markets. Each agent behaves alternatively according to its integrated producer, consumer, buyer, and seller classes. So each agent plays the four roles in turn. Agents are defined by several key numbers, which include the following.

- *Money, utility, stock:* Every agent starts with the same amount of money at the beginning of the simulation. This allotment of money allows the agent to produce goods, in which case the money leaks out of the system, or to purchase them from one of the other agents during a transaction. Agents also have an amount of accumulated utility that they increase by consuming products. The way the utility increases depends on the agent's tastes. Agents also possess a certain amount of products that they have produced but not yet sold. For logging purposes, these stocks are valued at current market prices. At the beginning of the simulation, agents start with no inventory.

- *Skills and tastes:* Agents are also characterized by two crucial vectors: their skills and their tastes, one of each corresponding to each product. This is the departing point of agents' differentiation during the simulation. Skills determine the cost of producing goods while tastes determine the amount of utility an agent will get from consuming a product. Skills and tastes are constant for each agent; once they are set, they will not evolve as the simulation proceeds. Furthermore, while individual skills and tastes vary, the total amount of skills ($\sum_i \text{skill}_i$) and tastes ($\sum_i \text{taste}_i$) is identical for every agent, hence agents only differ in the breakdown of their skills and tastes between products. At the initialization of the program, skills (tastes) are randomly set between 0 and 1 for each product; then they are all scaled so that the sum of all skills (tastes) for each agent is equal to 1.

- *Consumer behavior:* When an agent purchases a product, it is immediately consumed and converted into utility. Tastes determine the amount of utility that will be produced; for each unit of product p consumed, the agent's utility will increase by $\lceil \text{taste}_p \rceil$.

- *Producer behavior:* During each tick, an agent chosen randomly produces one unit of product. To produce a unit of product p, the agent will lose an amount of money determined by $\lceil \text{skill}_p \rceil$. This is the only process in the program that dissipates money (all the other processes lead to money transfers between agents and/or to utility increases).

2.3 The Agents' Choices

- *Producer choices:* A crucial part of the process is the selection of which product to produce. In a specification designated as *random production* (or zero-intelligence production), the product will be selected randomly, whereas in one designated as *informed production*, the agent will search for the product that is expected to maximize profit. Because production costs are known for each product, only the selling price needs to be estimated. This is done by querying the market for estimated prices, or in case none are available yet (which is the case in the first ticks before any transaction has taken place), by guessing at random.

Producers can therefore learn about the market supply and demand conditions. The market keeps a price estimate for every product, which is a moving average of the last *n* transaction prices for that product. Having an expected price and a given cost, the agent can calculate the expected profit (expectedPrice – cost) for each product and, knowing the product with the highest expected profit that can be affordably produced, select the kind of product to make.

- *Buyer and seller choices:* Buyers and sellers choose products on which to place offers based on their maximum expected profit (for sellers) or net utility (for buyers). To do so, sellers maximize the difference between their production costs and the latest price reported by the market for each product (see Section 2.5). Buyers make up their choice based on their expected utility (i.e., their taste for the product) and the latest price reported by the market for each product. However, in the random market structure, buyers—contrary to sellers—are offered one specific product to buy and can only accept or reject the offer based on their tastes. They will always accept an offer for a product at a price that is below their utility for the product. In both markets, buyers and sellers are both constrained: sellers only put offers on products they have in stock and never sell at a loss (i.e., their minimum selling price is their production cost), and buyers never place offers beyond their budget constraint and never put or accept offers at a price higher than their taste (utility).

2.4 The Ticks

The world moves forward through *ticks*, which are discrete time steps. During a tick, an agent is given a chance to produce one item; if the selected agent cannot produce anything, another agent is selected until one unit of a product has been produced. The market will then execute one transaction. As a result, one product unit is exchanged between the seller and the buyer agents. Once acquired, the buyer agent immediately consumes the product and the utility increases due to the agent's taste for the consumed product.

2.5 The Markets

Markets are the core of the model and represent the rules governing product exchanges between the agents. They also determine the buying and selling behavior the agents will play out. Two kinds of markets are studied in the model: the *competitive* one, which is a double auction market, and the *random/distributive* one.

- *Competitive market:* In the competitive market, agents bid to buy and sell goods. During a succession of steps, the market repeatedly invites two randomly selected agents to place asks and bids on one product they want to sell or purchase. This product selection can be done in different ways that are described in Section 3. In our baseline simulations, buyers choose the product that maximizes their expected net utility given the latest competing offer, and sellers choose the product that maximizes their expected net profit given the latest competing offer. In other settings, agents query the market to learn about actual prices, or use no market information but only base their choices upon their skills

or tastes, respectively. For each product, the market remembers only the highest bid and the lowest ask made in the current tick. As soon as these two numbers cross, the transaction is executed between the two winning agents for that specific product. If after a predetermined number of trials (arbitrarily high) no transaction can occur, the execution of the model stops and a market failure is reported. Agents place bids and offers in a manner inspired by Gode and Sunder [1]. Similarly to their work, our agents are faced with budget constraints, that is, they cannot spend more than the money they have, cannot buy a product at a price above their utility, and cannot sell a product below their production cost.

■ *Random/distributive market:* In the random market, consumers are proposed a certain product to buy from a given producer. Provided that the taste of the agent for the proposed product is larger than the producer's cost (inverse of skill) for the given product and that the buyer is sufficiently endowed, the transaction is made and the price is randomly set by the market between these two bounds. Buyers and sellers therefore do not learn anything from the market in the random model and are therefore closer to the zero-intelligence agents of Gode and Sunder. Here again, if following a predetermined number of trials, no transaction can occur, so a market failure is reported.

3. Results

3.1 Key Metrics

The first group of metrics quantifies the amount of welfare the agents accrued over time as a result of the transactions. There are two main dimensions of welfare: money and utility. Utility measures the amount of satisfaction the agents accumulated from their product consumption. It refers here to lifetime accumulated utility, so it is a monotonic increasing function, both at the aggregate and at the individual level. Money is the most obvious metric. It can be seen as a form of *potential utility* as it can be used to buy and consume products. Money leaves the world when agents produce and there is no way to inject more money back into the system (i.e., our model does not allow for any endogenous growth). Therefore, at the aggregate level, money is a monotonically decreasing function. However, at the individual agent's level, money can increase after a sale, though the overall trend will always be oriented downward. Our welfare comparisons therefore look at the amount of utility that is obtained by the consumers and the amount of money that is consumed by the production of goods. Total wealth is finally defined here as the sum of money and utility.

The second feature of the market to be examined is the amount of inequality it generates. The method used here to measure this inequality is the traditional Gini coefficient [11]. The Gini coefficient can be defined as twice the area between the Lorenz curve and the perfect equality line. As the data generated in this simulation gives a polygonal Lorenz curve, a simplified method [12] is used to calculate this

coefficient. It varies between 0 and 1 with 0 meaning perfect equality and 1 meaning perfectly inequality. So the closer to 1, the more unequal the market is. The Gini coefficient is used to measure inequality, not only for total wealth but also for money and utility separately.

3.2 Simulation Results

In our baseline simulations, the world was set with 50 agents, 10 products, an initial endowment of 500 money units per agent, and configured to keep a log of the 10 latest transactions for each product. Whatever the initial conditions we set in the model, the competitive market consistently and significantly produces more welfare at the aggregate level than the random market. This gain in welfare is mostly due to a gain in utility for consumers in the order of 60% overall. Given that no money is produced in our model but only consumed (by the production of goods) or transferred between buyers and sellers (with no transaction costs), no difference significantly appears in the amount of money that is left with the agents at the end of our simulations. This result, which confirms the earlier results of Gode and Sunder [1], is explained by a more efficient matching of consumers and products based on the preferences of the former. Each transaction therefore provides more utility to the consumer than what would statistically be achieved when the matching is purely random.

Nonetheless, our model allows for two possible modes of production choices by the producers: either *random* (agents choose which product to make essentially in a random way) or *informed* (agents choose which product to make based on the expected profit they could make with each product, itself computed based on their skills, which determine their costs of production, and the latest transaction prices observed on the market). When production choices are informed instead of random, considerably more money is left with the agents at the end of the simulations, indicating that the production is much more efficient (less money is wasted in producing goods for which the producer is less skilled). Overall, total cumulated production costs at the end of the simulations are 60% lower at the aggregate level when production choices are informed. In addition, when production choices are random, one simulation out of four is interrupted before the end of the 50 000 ticks due to the impossibility at some point to further match products with consumers' preferences, no matter the type of market. This result simply highlights the well-known benefits of specialization and comparative advantages theory. The occurrence of simulation failures in the presence of a random production is interesting computer-based evidence of Hayek's premonition of the market efficiency to exploit prices as an information system.

Figure 2 reports statistics over the key indicators of the state of the economy at the end of the simulations (200 distinct simulations were run sequentially and the statistics reported in the figure are averages over the 200 simulations). The figure reports results for the two types of markets (random and competitive) and for the two distinct methods of production choices (random and informed). Note that in each run, the four possible combinations of markets and production choices strategies were tested successively with the exact same sets of agents and products and the exact same initial conditions to ensure the comparability of the results.

Informed Production			
	Random	Competitive	% Difference
Total Utility	5857	9433	61%
Initial Endowment	25000	25000	
Total Money	23264	23242	0%
Total Money Spent	1736	1758	1%
Total Wealth	29122	32675	12%
Gini Utility	0.07	0.39	484%
Gini Money	0.01	0.10	638%
Gini Wealth	0.01	0.05	437%
Ticks per Run	50000	50000	
Runs	200	200	
Failure Rate	0%	0%	

Random Production			
	Random	Competitive	% Difference
Total Utility	7092	9927	40%
Initial Endowment	25000	25000	
Total Money	20758	20772	0%
Total Money Spent	4242	4228	0%
Total Wealth	27849	30698	10%
Gini Utility	0.13	0.41	212%
Gini Money	0.04	0.15	279%
Gini Wealth	0.01	0.03	243%
Ticks per Run	50000	50000	
Runs	200	200	
Failure Rate	24%	23%	

% Difference Across Production Methods (from informed to random)		
	Random	Competitive
Total Utility	-21%	-5%
Gini Utility	-93%	-3%
Gini Wealth	-9%	30%

% Change in Differential Performance of Models			
	Random	Informed	Difference
Total Utility	40%	61%	53%
Total Wealth	10%	12%	19%
Gini Utility	212%	484%	128%
Gini Money	279%	638%	129%
Gini Wealth	243%	437%	80%

Figure 2. Table with simulation results.

What is more appealing in our results is the huge differences that appear between the two types of markets in terms of welfare distribution. At the end of each simulation, we compute the concentration of utility, money, and their sum (representing the total wealth of the agents) across agents with a Gini index. A higher value of the index indicates that the utility, money, or wealth is more concentrated (i.e., more unequal).

Figure 2 reports these indexes for each of the four combinations of markets and production choice strategies. It clearly appears that the competitive market leads to much more inequality in the distribution of utility and money (and hence wealth) than does the random market. With informed production, utility ends up five times more concentrated in the competitive than in the random market, revealing considerable inequalities in the distribution of utility between agents. Like utility, money is much more concentrated in the competitive than in the random market (the Gini index is 6.5 times larger). Given that these differences are robust to a variety of changes in the model and in its initial conditions as well as to a large number of successive simulations, these results can be taken as proof that, all other things being equal, a competitive matching of consumers and producers based on price competition leads to much more inequality in welfare distribution than a purely random matching. In other words, a competitive market generates significantly more welfare but distributes it much more unequally between agents.

To help understand this phenomenon, recall that agents enjoy some utility from consuming products but they lose the money that corresponds to the price they had to pay for the products. Therefore, any transaction in the model distorts the distribution of utility and money by transferring some money from the consumer to the producer and some utility in the reverse direction. In the random market, consumers and producers are selected randomly. Therefore agents should statistically behave as producers and consumers with similar frequencies, so that transactions should statistically compensate for the distortions created by the previous ones, explaining why utility and money tend to be evenly distributed across agents in our random market. In contrast, in the competitive market, consumers and producers are matched based on their tastes and skills, respectively. This implies that products have a propensity to go more frequently to the consumers that have more differentiated preferences across products as their taste for a given product is larger than that of most other agents. Provided that they do not hit their budget constraint, they will therefore be willing to pay a higher price than the other consumers to get the same product and will therefore be favored in the matching process (i.e., they will win their auctions more often and will therefore statistically spend more money and accrue more utility). Likewise, producers with more pronounced skills will enjoy lower production costs than most other agents for a given product and will therefore be willing to offer consumers a cheaper price for their prod-

ucts. As a result, they will be favored in the auctions and will tend to accrue more money. Therefore, distortions in utility and money tend to grow over time. The competitive market favors those with skill in demand and those with taste skillfully satisfied. The self-amplifying pairing between skills and tastes underlines the growing inequality.
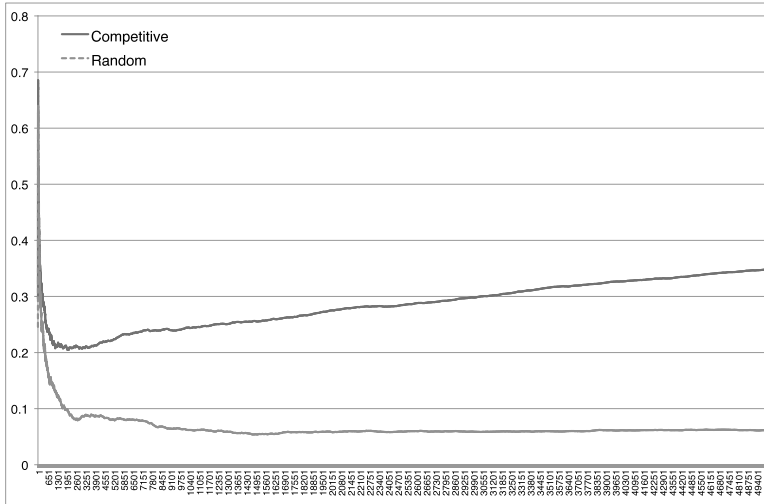

Figure 3. Evolution of Gini index for the random and the competitive markets.

This is illustrated in Figure 3, which exhibits the evolution of the Gini index of utility along the 50 000 ticks (iterations) in a simulation of the two markets with informed production. The very high value of the curve at the beginning of the simulation is artifactually due to the initial conditions made equal for all agents, so that the first transactions entail huge differences among agents. More significantly, the results show that the Gini index quickly stabilizes over 0.06 in the random market, whereas in the competitive market it starts at about 0.2 (once some initial transactions have taken place) and then increases almost linearly until the end of the simulation, indicating that inequalities continue to grow over time. The differences are in fact even more pronounced with informed than with random production choices. A closer look at Figure 2 shows that the Gini index of utility (money) is 50% (66%) smaller with informed than with random production in the random market, whereas Gini measures are less significantly impacted by production choice strategies in the competitive market. This suggests that all agents are better off with informed than with random production in a random market, probably because they all benefit from the advantage of specialization (i.e., goods are made by the most efficient producers and are therefore offered at a more affordable

price, making them attractive to agents with less pronounced preferences). In the competitive market, informed production only generates benefits in terms of money distribution, which ends up less concentrated (i.e., less unequal) with informed than with random production, probably due to more efficient production leading to cheaper prices, so that voracious consumers need to spend less money to satisfy their preferences. Consumers with more differentiated tastes across products will still be favored in the auctions and will therefore continue to accrue more utility than the others, but they will need to spend less money to do so and the distortions in money distribution will therefore be smaller with lower production costs. An informed production does not really increase the utility inequality in a competitive market; while more tasty products are at the disposal of competitive consumers, the lower price makes them also available to less greedy ones. All these results provide additional support for the positive impact of an informed production. Nevertheless, randomness has a positive impact in the market (in its random/distributive version, the pairing of the sellers and the buyers is purely random) when equality is more at stake than aggregate utility.

Finally, we investigated the impact of different behaviors (or degrees of learning) of the agents in the competitive market by testing three different criteria for their selection of products for which to place offers. In all the simulations we reported on, sellers and buyers make their sale and purchase decisions respectively based on the gap between the last bid that has been made over the same product by a competitor and their production costs or preferences. This means that they choose a product that would maximize their expected benefit (in terms of net profit or net utility) if they were to win the auction by slightly outbidding the best competing offer so far. We explored two other product selection rules in the competitive model. The first alternative rule is the same as our default rule, except that the agents base their estimations of expected profit on the last transaction prices for the product rather than on the latest competing bid. These two strategies already point to different learning processes by the agents. The second alternative rule simply lets the producers select the product at which they are best skilled and can afford to produce, and the consumers select the product that they have the stronger inclination (taste) for and can afford to purchase. In this latter model, agents therefore do not use any information from the market to choose which product to sell or buy.

Figure 4 reports the key average metrics over 200 sets of simulations with each of these three production selection rules, and with two different money endowment strategies: our default one (a 500 unique endowment at the beginning) as well as a scheme in which agents receive a very small amount of money at the beginning (2 units) and are granted a 0.002 unit at the end of each tick. All of the reported simulations are based on informed production choices. Although our main findings about the superior performance of the

competitive market over the random market in terms of efficiency and the higher inequality it leads to are rather robust, the magnitude of the differences between competitive and random markets vary from one simulation to another. This is particularly due to inequalities in utility in the competitive market that significantly depend on the product selection strategy used. This result confirms that, as suggested by Gode and Sunder, the behavior of the agents can influence the distribution of profits (but more importantly of utility) in the competitive market. Combined with our first sets of results, however, this indicates that institutions can generate inequalities (the competitive market always leads to more inequality than the random market design), but that the behavior of the agents can either reinforce or reduce these inequalities without much effect on the overall efficiency of the institutions in welfare creation.

	Baseline Latest Bid		Model 2 Skills/Tastes Only		Model 3 Latest Transaction		Model 4 Latest Bid		Model 5 Skills/Tastes Only		Model 6 Latest Transaction	
Product Selection												
Initial Endowment	500		500		500		2		2		2	
Money Increase	0		0		0		0.002		0.002		0.002	
	R	C	R	C	R	C	R	C	R	C	R	C
Utility	5857	9433	5863	8074	5865	7155	5877	9094	5878	8012	5875	7239
Money Spent	1736	1758	1756	1687	1732	1639	1745	1689	1756	1658	1772	1683
Gini Utility	7%	39%	7%	24%	7%	13%	7%	29%	7%	21%	7%	14%
Gini Money	1%	10%	1%	5%	1%	2%	9%	44%	9%	29%	10%	18%

Figure 4. Simulations with different endowments and production selection strategies.

A variety of tests has been performed to assess the robustness of our results to the main parameters of the model, such as the agents' skills and tastes. These are randomly allocated at the start of each set of simulations (i.e., we always simulate the two markets, random and competitive, with the exact same parameters and agents). By running multiple sets of simulations, we test the sensitivity of the results to the allocation of tastes and skills. We further checked whether our results are robust to a variety of combinations of parameters, including the money endowment of the agents, the number of agents, and the number of products. They are. In the interest of space, the results of these robustness tests are available from the authors upon request.

▎ 4. Conclusions

The objective of this paper is to examine to what extent a competitive market compares to a random market in welfare creation and above all in welfare distribution among agents. Although the economic literature usually attributes a higher efficiency to competitive markets in maximizing social welfare, very little attention has been paid so far to the equality in welfare distribution resulting from the competition between agents, notably due to the difficulty in solving such problems analytically with a large number of products and agents.

Agent-based simulations such as the one reported in this paper enable studying such emerging properties from individual interactions between agents. Although the random model could be handled in a more classical mathematical way due to the simplicity of the underlying rules, this would definitely not be the case of the competitive one. For obvious reasons and to ease the comparison, we maintain the same style of agent-based modeling for both models. Various conclusions may be drawn from our results.

First, while it creates more welfare (utility and money) at the aggregate level, the competitive market distributes it much less equally. The competitive market structure is responsible for an inequality amplifying effect: goods become concentrated in the hands of greedy consumers and money in the hands of skillful producers. By matching their tastes and skills, any pair of winning agents (the producer and the consumer) will benefit from the competitive nature of the market.

This result shows that institutions can in and by themselves generate inequality. Second, the behavior of the agents and the information they learn from the market (especially in choosing which products to bid on) have little effect on welfare creation (which is consistent with Gode and Sunder's results), but do significantly impact the distribution of welfare. This suggests that both institutions and agents share the responsibility for inequalities.

Various sources of randomness in real life are well known to compensate for the positive feedback resulting from competition. For instance, among others, competing agents have limited time and cognitive resources to explore all possible offers, and many apparently irrational motivations undermine a lot of trading decisions. While a lot of casualties make markets diverge from ideally competitive interaction in practice, how a fully random market could be practically designed in real life, although an interesting question, is out of the scope of this paper. However it is, for instance, quite plausible to imagine a computer-based market (such as eBay) where, after the seller has proposed a product and the lowest offer he is likely to accept, and the consumers, hidden from each other, indicate the offer they propose for that same product, a transaction takes place based on one possible random pairing.

Despite the care we took in testing the sensitivity of our findings to arbitrary choices in design and in initial conditions, these results will not allow us to generalize our findings to any market design. These results describe a stylized exercise in which we compare a very aggressive competition-based market mechanism in the form of a double auction with a pure theoretical abstraction that represents a market in which producer and consumer matching would be made purely on a random basis under a limited (and natural) set of constraints: the budget constraint, the "no sale at loss" rule for the producers, and the "no purchase above utility value" rule for the consumers. There are clearly large avenues for further analysis on how our results would change with other market mechanisms and agent behaviors.

References

[1] D. K. Gode and S. Sunder, "Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality," *The Journal of Political Economy*, **101**(1), 1993 pp. 119–137. doi:10.1086/261868.

[2] A. K. Sen, *On Economic Inequality*, Oxford: Oxford University Press, 1973.

[3] A. M. Okun, *Equality and Efficiency, the Big Tradeoff*, Washington, DC: The Brookings Institution, 1975.

[4] K. G. Binmore, *Natural Justice*, New York: Oxford University Press, 2005.

[5] L. Tesfatsion, "Agent-Based Computational Economics: Growing Economies from the Bottom Up," *Artificial Life*, **8**(1), 2002 pp. 55–82. doi:10.1162/106454602753694765.

[6] B. LeBaron, "Agent-Based Computational Finance," in *Handbook of Computational Economics*, Vol. 2 (L. Tesfatsion and K. Judd, eds.), North-Holland Publishing, 2006 pp. 1187–1233.

[7] B. LeBaron, "Agent-Based Computational Finance: Suggested Readings and Early Research," *Journal of Economic Dynamics and Control*, **24**(5–7), 2000 pp. 679–702. doi:10.1016/S0165-1889(99)00022-6.

[8] L. Tesfatsion, "Agents Come to Bits: Towards a Constructive Comprehensive Taxonomy of Economic Entities," *Journal of Economic Behavior & Organization*, **63**(2), 2007 pp. 333–346. doi:10.1016/j.jebo.2005.12.016.

[9] W. B. Arthur, J. H. Holland, B. Le Baron, R. G. Palmer, and P. Taylor, "Artificial Economic Life: A Simple Model of a Stockmarket," *Physica D: Nonlinear Phenomena*, **75**(1–3), 1994 pp. 264–274. doi:10.1016/0167-2789(94)90287-9.

[10] J. Cassidy, *How Markets Fail: The Logic of Economic Calamities*, New York: Farrar, Straus, and Giroux, 2009.

[11] C. Gini, "Measurement of Inequality of Incomes," *The Economic Journal*, **31**(121), 1921 pp. 124–126.

[12] B. Slack and J.-P. Rodrigue. "The Gini Coefficient." (Apr 13, 2012) http://people.hofstra.edu/geotrans/eng/ch4en/meth4en/ch4m1en.html.