

Two Tales of Complex System Analysis: MaxEnt and Agent-Based Modeling

Jangho Yang^{1,2} and Adrián Carro^{2,3}

¹*Oxford Martin Programme on Technological and Economic Change*

²*Institute for New Economic Thinking at the Oxford Martin School*

³*Mathematical Institute, University of Oxford*

Abstract

Over the recent four decades, agent-based modeling and maximum entropy modeling have provided some of the most notable contributions applying concepts from complexity science to a broad range of problems in economics. In this paper, we argue that these two seemingly unrelated approaches can actually complement each other, providing a powerful conceptual/empirical tool for the analysis of complex economic problems. The maximum entropy approach is particularly well suited for an analytically rigorous study of the qualitative properties of systems in quasi-equilibrium. Agent-based modeling, unconstrained by either equilibrium or analytical tractability considerations, can provide a richer picture of the system under study by allowing for a wider choice of behavioral assumptions. In order to demonstrate the complementarity of these approaches, we use here two simple economic models based on maximum entropy principles—a quantal response social interaction model and a market feedback model—for which we develop agent-based equivalent models. On the one hand, this allows us to highlight the potential of maximum entropy models for guiding the development of well-grounded, first-approximation agent-based models. On the other hand, we are also able to demonstrate the capabilities of agent-based models for tracking irreversible and out-of-equilibrium dynamics as well as for exploring the consequences of agent heterogeneity, thus fundamentally improving on the original maximum entropy model and potentially guiding its further extension.

Keywords: maximum entropy, agent-based model, entropy-constrained, irreversible process, heterogeneous agents

JEL Class: A10, D80, C63

1 Introduction

The past four decades have seen a growing number of contributions applying concepts from complexity science to a broad range of problems in economics. Motivated in part by dissatisfaction with conventional analytical approaches to economic problems, these contributions have applied, among others, Information Theory, Agent-Based Modeling, and Network Theory to offer a series of important methodological, theoretical, and empirical innovations. For example, the information-theoretic approach has gained a great amount of attention in macroeconomics since Christopher Sims' Rational Inattention (RI) program, in which the notion of mutual information is utilized in the macro model to predict the sluggish and erratic responses of economic agents (Sims, 2003, 2006). Agent-based models have also been widely utilized in various economic questions and have been recently included within the analytical toolbox of major policy institutions (Baptista et al., 2016).

The emergence of these new analytical approaches has highlighted the existence a series of methodological challenges within the standard economic theory. Notably, the conventional notion of equilibrium as a fixed state of the system, such as a market-clearing Walrasian equilibrium, has been called into question by these new approaches' understanding of equilibrium as an emergent property of the system that comes with endogenously generated fluctuations. In a similar vein, these new methods acknowledge the limitations of the rational choice theory in economics and understand human behavior as fundamentally constrained by uncertainty. Among many strands of complexity economics, information-theory-based approaches, especially Maximum Entropy (MaxEnt), and agent-based modeling have been prominent in raising these methodological issues with standard economics. Agent-based modeling is a simulation-based approach that traces the emergence of collective phenomena in a complex system back to the actions and interactions of its constituent agents. This method is highly flexible and can describe a wide range of complex system phenomena, e.g., the emergence of equilibrium, endogenous bifurcations, and transitions between multiple equilibria. Maximum entropy modeling is a quasi-equilibrium approach that analytically derives the equilibrium state of the complex system with a residual variation. Unlike the deterministic notion of equilibrium, MaxEnt predicts the central tendency of the system along with its endogenous fluctuation simultaneously.

This paper examines MaxEnt and agent-based modeling and argues that these two seemingly unrelated approaches can actually complement each other, providing a powerful conceptual/empirical tool for the analysis of complex economic problems. This is primarily because the weaknesses of one method can be overcome by the strengths of the other. For example, while agent-based models are highly flexible and can capture rich aspects of the economy as a complex system, their very flexibility leads to the issue of a high degree of freedom in model specification and thus makes them less tractable than standard methods. In contrast, while MaxEnt has a closed-form solution and thus helps us understand the qualitative properties of the model, it cannot easily accommodate as wide a choice of behavioral and institutional assumptions as agent-based models can. Benefiting from the strengths of each approach, we will show that MaxEnt can provide a testable analytical solution to quasi-equilibrium problems, from

which an agent-based version can be implemented to provide a richer picture of the model and guide its further extension. Especially, we will show that agent-based modeling can enable us to examine the out-of-equilibrium and possibly irreversible behavior (transitions between equilibria) of models based on MaxEnt considerations.

To demonstrate how MaxEnt and agent-based modeling can complement each other in understanding a complex economic system, we will study two simple models: a quantal response social interaction model and a market feedback model. In the first example, a typical agent tries to mimic other agents' decisions in deciding on her own action with uncertainty. As Foley (2017) shows, the MaxEnt solution to this type of problem in its simplest setting leads to two qualitatively different solutions: one with two stable extreme equilibria and one unstable internal equilibrium and the other one with a single stable equilibrium. We will show that an agent-based equivalent model can reproduce the outcomes of this model and can track the out-of-equilibrium dynamics of the system as well. The second example focuses on how the decisions of an economic agent have an impact on the market and how this has, in turn, an effect over the decisions of other agents. In particular, we focus on how firms' investment decisions impact the overall price/profit in the market. We will show that the MaxEnt solution leads to a 4-parameter probability distribution (Scharfenaker and Foley, 2017), whose general pattern can be reproduced and be further extended to a heterogeneous agent case within the agent-based equivalent model.

The rest of the paper is organized as follows: in Section 2 we briefly describe MaxEnt and agent-based modeling, focusing on their economic applications. In Sections 3 and 4, we compare both approaches by using two economic models based on MaxEnt considerations and deriving agent-based equivalent models. In particular, we focus, respectively, on a quantal response social interaction model and a quantal response statistical equilibrium model. Section 5 draws some conclusions.

2 Entropy Constrained Modeling and Agent-Based Modeling

MaxEnt: Entropy Constrained Economic Model

Shannon entropy $H[p]$, defined as the negative expected value of the log-probability, $-\sum_x p[x] \text{Log}[p[x]]$, is the correct form of the average *information content* of a system composed of x states (Shannon, 1948; Jaynes, 2003; Cover and Thomas, 2006; MacKay, 2005) with the convention $0 \text{Log}(0) = 0$. Information content, $I(p) = \log[1/p]$, is monotonically decreasing in p , which implies that the more unlikely the outcome is, the more information content the outcome provides. In other words, the occurrence of a highly likely event (high p) does not have a high degree of information content while a highly unlikely event (low p) provides much information. This notion of information content (or sometimes called “surprisal”) can directly translate to the notion of uncertainty. When we are less certain about the system (low p), we expect to get more information from observing the actual occurrence of a set of outcomes: the higher the uncertainty, the more information content in the system. Therefore, the Shannon entropy,

which is the average information content of the probability distribution, can be interpreted as a degree of uncertainty. The higher the entropy, the more uncertain we are about the system, and the more information content we can obtain from observing the actual occurrence of an outcome. For a detailed discussion on the various aspects of entropy in physics and information theory, see Bais and Farmer (2007).

There have been two different ways of using the concept of entropy in economic applications (Yang, 2018). The first line of models maximizes the entropy of a target economic variable —e.g., the returns on capital (Scharfenaker and Foley, 2017; Scharfenaker and Semieniuk, 2017), firms' growth (Irle et al., 2008), Tobin's q (Scharfenaker and dos Santos, 2015)— subject to constraints based on some relevant theory. The second line of models maximizes the expected payoff of individual economic agents subject to the entropy constraint (Sims, 2003; Foley, 2017). In this section, we will illustrate the maximum entropy approach in the context of the second line of entropy models, which are more akin to the standard economic models.

Consider a mixed strategy problem in which the typical agent is a payoff maximizer and is subject to a quantal decision (a decision on two discrete choices), $a \in \{0, 1\}$. The agent assigns a probability to each action, $P(a)$, and tries to maximize the expected payoff given the payoff function $U(a)$ (von Neumann and Morgenstern, 1944),

$$\begin{aligned} \max \quad & \sum U(a)P(a), \\ \text{s. t.} \quad & \sum P(a) = 1. \end{aligned} \tag{1}$$

With no further constraint, the solution to this problem is the Dirac Delta function δ ,

$$P^*(a) = \delta[a - \hat{a}[u]], \tag{2}$$

where $\hat{a}[u]$ is the payoff-maximizing behavior. This implies that the equilibrium is either $a = 0$ or $a = 1$ depending on the payoff structure.

Now suppose that agents are exposed to a positive degree of uncertainty due to the unobserved factors that prevent the agent's choice of the first-best outcome. To put it differently, agents are informationally constrained and have a limited capacity to process market signals. Two notable strands of information-constrained models are Sims' Rational Inattention model (Sims, 2003) and Foley's entropy constrained model (Foley, 2017). First, Sims (2003) proposed the Rational Inattention model, in which optimizing agents make economic decisions by processing external random signals with a constraint of finite channel capacity that determines the maximum information flow. Since the channel capacity is defined as the maximum of the mutual information between the action and signals, Sims' model is a mutual information constrained model.¹ Second, Scharfenaker and Foley (2017) and Foley (2017) proposed a more simplified version of an entropy constrained behavioral model, in which payoff maximizing agents

¹The mutual information of an input and an output $I(X, Y)$ is defined as the change in information after observing the

are directly constrained by the Shannon entropy, not by the mutual information. The agent's optimal strategy in the entropy constrained model results in choosing probabilistically in line with a logistic quantal response behavior, which is essentially equivalent to Sims' theory of rational inattention (Foley, 2017).² The paper follows the second approach by Scharfenaker and Foley (2017). For a more detailed comparison of these two approaches, see Yang (2018).

Extending the simple mixed-strategy model to the entropy-constraint model, we have the following optimization program with an additional minimum entropy constraint

$$\max \quad \sum U(a)P(a), \quad (3)$$

$$\text{s. t.} \quad \sum P(a) = 1, \quad (4)$$

$$- \sum P(a) \log[P(a)] \geq H_{min}, \quad (5)$$

where H_{min} represents the minimum entropy of the typical agent. The standard method to solve the constrained entropy program is to use the Lagrange duality (Borwein and Lewis, 1991; Cover and Thomas, 2006) as

$$L(P(a), \mu, T) = \sum U(a)P(a) - \mu \left(\sum P(a) - 1 \right) - T \left(- \sum P(a) \log[P(a)] - H_{min} \right), \quad (6)$$

where μ and T are the Lagrange multiplier (dual variables) associated with the primal problem. Note that the Lagrange multiplier T corresponds to temperature variables in statistical mechanical models. Following Scharfenaker and Foley (2017), we call this agent's behavior temperature. Since the objective function is linear in the frequencies and the entropy is a concave function, the first order conditions of the Lagrangian function are necessary and sufficient conditions to attain the global maximum. Note that the inequality constraint is treated in the same way as the equality constraints because the strong duality holds for entropy maximization problem (Boyd and Vandenberghe, 2004). Taking the first derivative

output Y given the prior information on X :

$$I(X; Y) = H(X) - H(X|Y) \geq 0,$$

where $H(X|Y)$ is a conditional entropy defined as $-\sum p(x, y) \log \left[\frac{p(x, y)}{p(y)} \right]$.

²When using the Shannon entropy, there exists a duality between the utility maximization with a minimum entropy constraint and the entropy maximization with a minimum utility constraint since both share the same Lagrangian function (Scharfenaker and Foley, 2017). This type of duality is well known in the entropy literature. For example, the entropy maximization problem with linear constraints for the multinomial logit model has the same solution as the likelihood maximization problem due to the same Kuhn-Tucker conditions in both problems (Anas, 1983). This makes the interpretation of the entropy-constrained model dual as well. The utility maximization with the minimum entropy constraint implies that the agents are assumed to be exposed to uncertainty in their decision-making process (bounded rational model), while the entropy maximization with the minimum utility constraint implies that researchers are exposed to uncertainty in predicting the agent's behavior according to their model (See Blackwell (2018) for a detailed discussion on the relationship between the logistic regression and the maximum entropy approach.).

with respect to $P(a)$, μ , T and setting to zero, we get

$$\frac{\partial L}{\partial P} = U(a) - \mu + T(1 + \text{Log}[P(a)]) = 0, \quad (7)$$

$$\frac{\partial L}{\partial \mu} = \sum P(a) - 1 = 0, \quad (8)$$

$$\frac{\partial L}{\partial T} = - \sum P(a) \log[P(a)] - H_{min} = 0, \quad (9)$$

whose resulting frequencies of a due to the equality constraint (8) are

$$P(a) = Z(U, T, a)^{-1} e^{-\frac{U(a)}{T}}, \quad (10)$$

where $Z(\cdot)$ is the normalizing partition function, $\sum e^{-\frac{U(a)}{T}}$. We call this resulting decision probability a *quantal response function*, which has been extensively used in the game theory literature (McKelvey and Palfrey, 1996; McKelvey and Palfrey, 1998). One of the differences between our entropy constrained model and the standard game-theoretic approach is that the former derives the quantal response function as an implication of the minimum entropy constraint, while the latter derives it as a result of a fixed point theorem. This quantal response function will be the basis of more complicated MaxEnt models in the following sections.

We now introduce variables affecting the actions of the agents. Suppose that the investor's decision to invest in the market is affected by some social variable x , e.g. the average returns on capital or other investors' investment strategy. This can be represented by the conditional frequency distribution $P(a|x)$, which we interpret that the binary decision $a = \{0, 1\}$ is dependent on some social variable x that is not under the direct control of the agent. Note that the payoff function now includes the social variable x as well since the agent's payoff is determined by how much the agent responds to the social variable. The optimization program is now written as

$$\begin{aligned} \max \quad & \sum U(a, x)P(a|x), \\ \text{s. t.} \quad & \sum P(a|x) = 1, \\ & - \sum P(a|x) \log[P(a|x)] \geq H_{min}. \end{aligned} \quad (11)$$

Program 11 has the same mathematical structure as the simple mixed-strategy model to the entropy-constraint model and solving in the same fashion leads to the following solutions

$$\begin{aligned} P^*(a = 0|x) &= \frac{e^{\frac{U(a=0,x)}{T}}}{e^{\frac{U(a=0,x)}{T}} + e^{\frac{U(a=1,x)}{T}}}, \\ P^*(a = 1|x) &= \frac{e^{\frac{U(a=1,x)}{T}}}{e^{\frac{U(a=0,x)}{T}} + e^{\frac{U(a=1,x)}{T}}}, \end{aligned}$$

where $P^*(a = 0|x)$ and $P^*(a = 1|x)$ are the maximum entropy probability of a binary action given social variable x . For example, suppose the social variable is the returns on capital (rate of profit) and the action variable is the binary entry-exit decision in the market. The typical firm i has a hurdle rate of profit μ_x , which denotes its desirable level of the rate of profit. If the observed rate of profit is greater than the hurdle rate, $x > \mu_x$, the firm is likely to enter the market. By the same token, when $x < \mu_x$, the firm is likely to exit the market. We suppose that the utility function is a linear function of x with μ_x as a reference, $U(a = 1, x) = x - \mu_x$. Therefore, the quantal response function of entry is $P(a) = 1/(1 + \exp(-(x - \mu_x)/T))$, where T represents the degree of uncertainty. Figure 1 visualizes this conditional quantal response function for four different levels of the behavior temperature.




Figure 1: A quantal response function with the linear payoff specification. The x-axis is the social variable as the conditional variable of the quantal response function, while the y-axis is the probability of action, $a = 1$, as a function of the social variable. Four different levels of behavior temperature are displayed in different colors. As the behavior temperature becomes larger, the quantal response function becomes more uniform. Except for $T = 0$, the quantal response function is non-degenerate.

The x-axis is the social variable as the conditional variable of the quantal response function, while the y-axis is the probability of action, $a = 1$. The higher T is the more uncertain the quantal decision is and thus the action probability is more spread out. When T is sufficiently large, the quantal response is almost uniform, implying that the agent does not respond to the social variable. Except for the unattainable case when $T = 0$, the quantal response function is non-degenerate and assigns positive probabilities to all possible states.

The quantal response function we derive in this section will be the basis of the two examples we will discuss in the following sections.

Agent-based modeling

Agent-based models are computer simulations consisting of a dynamic system of autonomous, heterogeneous, interacting agents. These agents are autonomous in the sense that each of them individually assesses its situation—its environment—and makes decisions according to a predefined set of behavioral rules. Different levels of heterogeneity are possible, from agents of the same type but different parameter values or (stochastic) histories to completely different types of agents altogether. The interactions between these agents, which can be mediated by different network structures, determine how the state and actions of one agent affect the state and actions of the others. In this way, cascades of individual changes of state are possible following a change in the environment, leading, at a system level, to a nonlinear emergent response. As opposed to traditional economic models, agent-based models are not solved for equilibrium, but rather simulated. This amounts to instantiating the population of agents and letting them assess their environment, interact and make their decisions for a number of time steps while monitoring all micro and macro variables of interest. The main goal of this approach is thus to explore how complex, collective, macro patterns in the economy emerge out of micro-level behaviors, interactions, and structures.

Following Axtell (2000), we can classify the uses of agent-based models into three broad categories. The first refers to settings in which the question under study can be completely formulated in terms of equations that are explicitly solvable, whether analytically or numerically. While an agent-based model is hence not necessary in these cases, it can be a useful tool for presenting the results of analytically solvable systems and it will function as a certain type of Monte Carlo simulation for systems that can be solved only numerically. The second category is formed by settings where equations can be written down but not completely solved, in which case, agent-based modeling can provide important insights about the general structure of the solution as well as the dynamical properties of the system, including any out-of-equilibrium behavior. Finally, the third kind of use of agent-based models is to explore in a systematic manner problems for which a mathematical formulation in terms of an explicit system of equations is not even useful, the complexity of this formulation not allowing for any advance in that way.

The contributions of this paper span across these three categories. In particular, we first reproduce well-known analytical results of two simple economic models based on maximum entropy theory. We then show how agent-based modeling can be used to explore the dynamics of these models, for which no analytical solution exists. Notably, for one of the models, this approach allows us to study the behavior of the system out of equilibrium, as we are able to observe the transitions between its equilibria, and thus build, for instance, statistics about the probability of observing a crossing between these equilibria and how this probability depends on the parameters of the model. Finally, we are able to expand on one of the models by introducing heterogeneity in the form of two types of agents with different parameter values, an extension which, due to its difficulty, has not yet been addressed within the available maximum entropy formalism.

3 Example 1: Social Interaction in the Market

MaxEnt: Quantal Response Social Interaction Model

As a first example, let us study the case where the social variable x in the conditional quantal response function is the average behavior of the other agents, \bar{a} . This is a situation where bounded rational agents respond to other agents' behavior (See Cherkashin et al. (2009) for a general social interaction model.). In the binary action case, $a = \{0, 1\}$, the the average behavior of the other agents is measured by the average frequency of action, $\bar{a} = \sum_{i=1}^N a_i / N$ where N is the total number of agents. The entropy-constrained optimization program given the average frequency of action is the following (Foley, 2017)

$$\begin{aligned} \max \quad & \sum U(a, \bar{a})P(a|\bar{a}), \\ \text{s. t.} \quad & \sum P(a|\bar{a}) = 1, \\ & - \sum P(a|\bar{a}) \log[P(a|\bar{a})] \geq H_{min}, \end{aligned} \tag{12}$$

whose solution for $a = 1$ is

$$P^*(a = 1|\bar{a}) = \frac{e^{\frac{U(a=1, \bar{a})}{T}}}{e^{\frac{U(a=0, \bar{a})}{T}} + e^{\frac{U(a=1, \bar{a})}{T}}} = \frac{1}{1 + e^{-\frac{U(a=1, \bar{a}) - U(a=0, \bar{a})}{T}}}.$$

The social interaction type of behavior implies that the probability of the typical agent's action is positively related to the average frequency of the action in the economy. For simplicity, suppose the payoff difference between the action and non-action is modeled as a linear function of \bar{a} ,

$$U(a = 1, \bar{a}) - U(a = 0, \bar{a}) = \bar{a} - \mu_a, \tag{13}$$

where μ_a is a reference probability of action, which can be set to 0.5 for convenience, meaning that the agent is indifferent between the two binary decisions $a = \{0, 1\}$ when exactly a half of other agents decide to act, $\bar{a} = 0.5$. As a result, the conditional quantal response function in the social interaction model becomes

$$P^*(a = 1|\bar{a}) = \frac{1}{1 + e^{-\frac{\bar{a} - 0.5}{T}}}. \tag{14}$$




Figure 2: Social interaction with a quantal response. The x-axis is the average frequency of the action while the y-axis is the individual action frequency given the average frequency. The black line is the equilibrium line where the individual agent behavior and the average behavior match. The red line is the quantal response function with a high behavior temperature T , while the green line corresponds to a low behavior temperature. When T is high there is a unique internal stable equilibrium. When T is low, there are two stable extreme equilibria and one unstable internal equilibrium.

Figure 2 visualizes this derived conditional quantal response function. The x-axis is the average frequency of the action while the y-axis is the individual action frequency given the average frequency. The black line is the 45-degree line and represents the equilibrium line where the individual agent behavior and the average behavior match. When there is lower uncertainty (lower T), as in the green line, there are two stable extreme equilibria where the decision probability is close to either 0 or 1 and one unstable interior equilibrium where the decision probability is 0.5. When the degree of uncertainty increases, as in the red line, there is a unique stable equilibrium. Note that this particular application of the MaxEnt on a binary decision with a social interaction (the second-order effect) is equivalent to Ising models (See for instance Xu et al. (2017); Castellana and Bialek (2014); Sakellariou et al. (2017)).

Agent-Based Simulation

In order to complement the MaxEnt results presented in the previous subsection, we develop here an agent-based model based on exactly the same rules of behavior assumed for the MaxEnt model. As we will show, this agent-based equivalent model will enable us to move beyond the MaxEnt analysis by allowing us to characterize the time evolution of the system as well as its out-of-equilibrium behavior.

In particular, let the system be composed by $N = 1000$ investors, each of them characterized by

a binary variable $a_i = \{0, 1\}$ setting their entry ($a_i = 1$) or exit ($a_i = 0$) from the market, and let $\bar{a}_t = \sum_{i=1}^N a_i / N$ be the average action or fraction of agents investing in the given asset at time t . Agent-based models are dynamic, and thus we need to specify the time evolution of the system. Furthermore, due to their algorithmic character, it is useful to specify this time evolution as a procedure in time, i.e., as a set of actions and updates that take place every time step. In order to start the dynamics we need to set exogenously an initial average action or initial fraction of agents investing, \bar{a}_0 , such that agents can use it to make their first time step decisions. Then, using the same decision rule as in equation (14) with $\bar{a} = \bar{a}_0$, as well as the values of the parameters μ and T , each of the N investors can find its initial probability of investing. Then a random number is uniformly drawn between 0 and 1 for each investor, who will decide to invest if this number is below the corresponding probability. Once the new market positions are computed, an updated value of the fraction of agents investing or average action can be found. Finally, this process can be repeated for a number of time steps: new probabilities of investing can be computed for the agents taking into account the previous average action, new random numbers drawn and thus new market positions set, and finally a new average action found.

Figure 3 shows the time evolution of the average action for four different levels of the behavior temperature T . As sketched by Foley (2017) and as described in the previous subsection, a transition can be observed from two equilibria to a single equilibrium as the temperature increases above a certain threshold or critical point, $T_c \sim 0.2495$. When the temperature is low ($T < T_c$), i.e., when agents have low levels of uncertainty, as in the top two plots ($T = 0.2$ and $T = 0.24$), two different equilibrium states can be clearly identified. For very low temperatures these equilibria are close to 0 and 1, fluctuations around them are small and no crossing between them can be observed within the plotted time window (see $T = 0.2$ in the upper left plot). As the temperature increases, i.e., as the level of uncertainty of the agents increases, the two equilibria move progressively away from 0 and 1 and closer to each other — and thus also closer to $\bar{a} = 0.5$ —, fluctuations around them become larger and more and more crossings can be observed between them (see $T = 0.24$ in the upper right plot). When the temperature increases above its critical value ($T > T_c$), the two equilibria merge into a single one —located at $\bar{a} = 0.5$ — and the system fluctuates around this unique equilibrium, as shown in the bottom two plots ($T = 0.25$ and $T = 0.3$). As can be observed by comparing these two plots, fluctuations around that single equilibrium become smaller the higher the temperature is above its critical value.




Figure 3: Average action or fraction of agents investing, \bar{a} , as a function of time, t , for four different behavior temperatures, T : well below the critical point ($T = 0.2$, with two colors identifying two different realizations of the process temporarily trapped in different equilibria), a bit below the critical point ($T = 0.24$), slightly above the critical point ($T = 0.25$), and well above the critical point ($T = 0.3$). Note that the critical point is at $T_c \sim 0.2495$.

The observed transition and the general behavior of the system can be explained in terms of the competition between the two mechanisms included in the model: a copy mechanism, by which agents copy what the majority of the other agents are doing, and a noise mechanism, by which agents are allowed to do otherwise, i.e., to take the minority action. While the copy mechanism tends to create large majorities of agents taking the same action, whether investing or not, the noise mechanism tends to hinder the emergence such majorities and to bring the system back to a random allocation of market positions. The behavior temperature parameter T regulates the relative strength of the noise mechanism over the copy mechanism, in such a way that a zero temperature implies that there is only copy and an infinite temperature implies that there is only noise. For temperatures above the critical point ($T > T_c$), the strength of the copy mechanism is not enough for stable majorities to emerge, and so the system fluctuates around a single equilibrium characterized by a random allocation of states, i.e., by half of the agents deciding to invest and the other half deciding to not invest. In this case, being the fluctuations

mainly driven by the copy mechanism trying to create majorities and thus move the system away from its equilibrium, their size decreases when the temperature increases (compare the lower left plot with the lower right plot in Figure 3). On the contrary, for temperatures below the critical point ($T < T_c$), the strength of the copy mechanism is enough for stable majorities to emerge, and so the system fluctuates around one of two possible equilibria: one with a majority of investors and the other one with a majority of non-investors. In this case, being the fluctuations mostly driven by the noise mechanism tending to destroy any majority and thus move the system away from the prevailing majority and towards a random distribution of states, their size increases when the temperature increases (compare the upper left plot with the upper right plot in Figure 3). Since it is precisely these fluctuations which allow the system to cross from one equilibrium to the other one, the likelihood of observing such crossings will also increase with temperature. Note that, in the limit case of a zero temperature, there are no fluctuations and thus no crossings between the equilibria: they become absorbing configurations and the system becomes non-ergodic.

For the sake of studying these fluctuations and crossings in a more quantitative manner, it is useful to define, for temperatures below the critical point ($T < T_c$), the equilibrium the system is at a given time step as the one corresponding to the equilibrium line most recently crossed by the trajectory of the system³. To this end, we need to find the position of both equilibrium lines, which we can achieve by numerically solving equation (14) in the case in which the probability of investing is equal to the actual fraction of investors,

$$\frac{1}{1 + e^{-\frac{\bar{a}-0.5}{T}}} = \bar{a}. \quad (15)$$

In this way, for example, we can identify 5 crossings between both equilibria⁴ in the 10000 time steps shown in the top-right panel of Figure 3. For the sake of continuity in the measures to be developed below, it is useful to consider, by extension, that there is a crossing also each time the system crosses the single equilibrium line at $\bar{a} = 0.5$ in the case of temperatures at or above the critical point ($T \geq T_c$).

In order to quantify the fluctuations, as well as their dependence on the behavior temperature, Figure 4 shows two different measures of the variance of the time series of the average action as a function of this temperature. The plot on the left shows the overall, unconditional variance of the fluctuations, i.e., treating the whole time series as a single trajectory. As a consequence, both fluctuations around a given equilibrium and fluctuations due to crossings between equilibria are aggregated under a single variance measure. On the contrary, the plot on the right shows the weighted average of two conditional variances: for temperatures below the critical point, these variances are conditioned on the system being in each to the two possible equilibria; for temperatures at or above the critical point, these variances are conditioned on the system being above or below the single equilibrium line. By independently considering fluctuations around each of the equilibria and only then aggregating them, this measure

³Note that this definition forces us to disregard the initial time steps, before the system crosses its first equilibrium line.

⁴Since the proposed definition is arbitrary, alternative definitions are possible, such as defining the equilibrium the system is at as the closest one to the current point in the trajectory of the system. Note, however, that this alternative definition would lead to identifying 15 crossings in the top-right panel of Figure 3, instead of the 5 identified by our method.

neglects the variance due to the system having two different equilibria, though it includes the variance due to the excursions from one equilibrium to the other.




Figure 4: Two measures of the variance of the time series of the average action as functions of the behavior temperature T . The plot on the left shows the overall, unconditional variance of the time series, $\sigma_{uncond.}^2(\bar{a})$. The plot on the right shows the weighted average of two conditional variances, $\sigma_{cond.}^2(\bar{a})$. For $T < T_c$, these variances are conditioned on the system being in each to the two possible equilibria; for $T \geq T_c$, these variances are conditioned on the system being above or below the single equilibrium line at $\bar{a} = 0.5$. Note that this second variance measure neglects the variance due to the system having two different equilibria, though it includes the variance due to the excursions from one equilibrium to the other.

Predictably, the unconditional variance (see plot on the left of Figure 4) increases as the behavior temperature becomes smaller due to the increasing distance between the two equilibria, i.e., driven by crossings between these equilibria. As the behavior temperature approaches zero, the unconditional variance converges to the expected value for a time series composed of zeros and ones, $\frac{1}{4}$. For very high temperatures, on the contrary, the unconditional variance tends, as expected, to the variance of a binomial distribution with success probability 0.5 and N trials (divided by N^2 since we are dealing with an intensive variable), $\frac{1}{4N}$. Note the somewhat counter-intuitive implication of this: the higher the level of individual uncertainty (higher temperature), the lower the level of system-wide uncertainty (lower variance); or, in the opposite sense, the more the agents copy the average behavior (lower temperature), the higher the uncertainty about the macroscopic state of the system (higher variance). Indeed, if each agent makes a random decision, then the system is very likely to be split in more or less equal parts, while, if each agent tries to coordinate with the other agents, then a majority will form for one or the other options, the specific option being unpredictable.

Regarding our alternative, conditional measure of variance (see plot on the right of Figure 4), defined so as to neglect the variance due to the existence of multiple equilibria, we first notice that it tends to zero as the temperature also tends to zero, coherent with the smaller fluctuations around equilibria that we observed in Figure 3 for decreasing temperatures below the critical point (compare upper plots in that figure). Second, we can observe that this conditional variance has a maximum right before

the critical point: as temperature increases below the critical point, the two equilibria move closer and closer, exerting a stronger and stronger pull on each other and thus allowing for larger and larger excursions of the system between them. Finally, for increasing temperatures above the critical point, the conditional variance decreases towards a finite non-zero value, equivalent, again, to a time series drawn from a binomial distribution with success probability 0.5 and N trials.

While any non-zero temperature below the critical point ($0 < T < T_c$) implies a non-zero probability of crossings between the equilibria, these crossings have a certain degree of irreversibility, i.e., once the system makes a crossing, it is likely to stay temporarily trapped around the new equilibrium. In any case, these crossings are never strictly irreversible, as the system will always return to the previous equilibrium given a long enough time window⁵. In order to characterize these crossings, let us define the probability per unit time of observing a crossing between the two equilibria of the system as the inverse of the average time the system spends in a given equilibrium before crossing over and reaching the other one, i.e., the inverse of the average time between two crossings, or, in other words, the frequency of crossings. This probability is shown in Figure 5 as a function of the behavior temperature. Note that, by extension, we also consider in this figure temperature values above the critical point, for which the frequency of crossings simply refers to consecutive crossings of the single equilibrium line at $\bar{a} = 0.5$. Given that this single equilibrium line is in fact the result of the merger of the two equilibria characteristic of the low temperature regime, this extension ensures continuity of the probability measure.




Figure 5: Frequency of crossings as a function of the behavior temperature T . For $T < T_c$, these crossings are between the two equilibria of the system; for $T \geq T_c$, these are crossings of the single equilibrium line at $\bar{a} = 0.5$. In both cases, this is a measure of the probability of observing such a crossing per unit time.

As shown in Figure 5, the probability of observing a transition between the two equilibria of the system is a monotonically increasing function of the behavior temperature. Interestingly, the slope

⁵Note that for $T = 0$ there are no crossings, and so the system stays forever trapped into the first equilibrium it reaches. In this sense, the system becomes strictly irreversible.

of this increase is characterized by significant changes: there is a first phase of exponential increase (note the logarithmic scale in the y axis) up to $T \sim 0.235$, then a phase of sub-exponential increase (note the decreasing slope), then an upwards bump around the critical point ($T_c \sim 0.2495$), and finally another phase of sub-exponential increase (note again the decreasing slope). Note that the first phase of exponential behavior implies that, for low and decreasing temperatures, the probability of observing a crossing vanishes exponentially for $T \rightarrow 0$.

In this example, we have developed an agent-based model with the same behavioral rules as a MaxEnt model in order to show how the former can complement and expand on the results of the latter. In particular, we have focused on a system with a transition between a phase with two equilibria and a phase with a single equilibrium. First, we have used the explicitly dynamic character of the agent-based model to explore individual trajectories or temporal realizations of the process. Second, we have characterized the fluctuations in these trajectories, as well as identified their main drivers, by considering two different variance measures. Finally, we have shown how agent-based modeling can be used to explore and characterize the behavior of the system out of equilibrium, when it is performing a transition between its equilibria. All of this is out of reach for the MaxEnt model.

4 Example 2: Actions and Market Feedback

MaxEnt: Quantal Response Statistical Equilibrium

We study the case where the economic agent's binary decision, $a = \{0, 1\}$ impacts the social variable $x \in \mathbb{R}$, e.g., a firm's investment decision impacts the overall price/profit in the market. This model adds another conditional frequency distribution $P(x|a)$ to the quantal response model in 3, which we interpret as the impact of the agent's decision a on the social variable x . Depending on the moment constraint on x given a , we can express a wide range of market feedback mechanisms in the MaxEnt setting. For a simple negative feedback case, which we will discuss in this section, the constraint takes the following form (Scharfenaker and Foley, 2017)

$$E(x|a = 1)P(a = 1) - E(x|a = 0)P(a = 0) = \int P(a = 1, x)x dx - \int P(a = 0, x)x dx = \xi. \quad (16)$$

This constraint compares the expected values of the social variable x when the agent decides to act ($a = 1$) and decides not to act ($a = 0$). For example, suppose that the social variable x is the profit rate and the action variable a is the investment decision on whether to invest or not. The above constraint compares the profit rates of the market when the investor has invested and has not. A high ξ means that the implicit market feedback mechanism is so weak that the profit rate after investors have already entered the market is still higher than the hypothetical situation with no investment. In contrast, when ξ is low, the implicit market feedback mechanism is so strong that the initially high-profit rate that induces investment is exhausted after the investors' entry decision, leading to a very small difference in the resulting profit rate between entry and exit.

The maximum entropy program with the quantal response action and the feedback effect can be written as

$$\begin{aligned}
\max \quad & - \int \sum_a P(a, x) \text{Log}[P(a, x)] dx, \\
\text{s. t.} \quad & \int \sum_a P(a, x) dx = 1, \\
& \int P(x) x dx = \psi, \\
& \int P(a = 1, x) x dx - \int P(a = 0, x) x dx = \xi, \\
& P(a = 1|x) = \frac{1}{1 + e^{-\frac{x-\mu}{T}}},
\end{aligned} \tag{17}$$

where we have also added a mean constraint on the social variable, implying that there is some reference level of the target social variable, e.g., the economy-wide rate of profit⁶. The maximum entropy distribution $P^*(x)$ is (Scharfenaker and Foley, 2017)

$$P^*(x) = \frac{e^{H_{x,\mu,T}} e^{-\beta \text{Tanh}\left[\frac{x-\mu}{2T}\right] x} e^{-\kappa x}}{\sum_x e^{H_{x,\mu,T}} e^{-\beta \text{Tanh}\left[\frac{x-\mu}{2T}\right] x} e^{-\kappa x}}, \tag{18}$$

where

$$H_{x,\mu,T} = H\left(\frac{1}{e^{-\frac{x-\mu}{T}} + 1}, \frac{1}{e^{\frac{x-\mu}{T}} + 1}\right), \tag{19}$$

being $H(x, y)$ the joint entropy of x and y , defined as $H(x, y) = - \sum_x \sum_y P(x, y) \log[P(x, y)]$, and where $\text{Tanh}[\alpha]$ is the hyperbolic tangent function and can be written as $\frac{e^{2\alpha}-1}{e^{2\alpha}+1}$. β is the Lagrangian multiplier of the ξ constraint while κ is the Lagrangian multiplier of the mean constraint. We call the resulting distribution in equation (18) the *Quantal Response Statistical Equilibrium* distribution (QRSE). Figure 6 shows a graphical characterization of the QRSE distribution, the marginal distributions of the social variable x for different parameter values.

⁶See Blackwell (2018) for a detailed discussion on various specifications of QRSE model.




Figure 6: QRSE distributions with different parameter values. μ , T , β , and κ parameters determine the location, the tail behavior, the scale and the skewness, respectively.

Note that μ and β determine the location and the scale, respectively, while T and κ determine the tail behavior and skewness, respectively. Predictably, μ , as the hurdle rate of the decision, shifts the distribution. T and β determine the dispersion of the distribution. A lower T means a more peaked distribution and thus a more occurrence of the extreme event. Note that the parameter β is inversely related with ξ . Therefore, smaller β means a higher ξ and thus a weak market response, making the distribution more dispersed. Finally, the higher κ , the more left-skewed the distribution.

Agent-Based Simulation

In order to complement the MaxEnt results presented in the previous subsection, let us now develop an agent-based model based on the same rules of behavior assumed in the MaxEnt model. In this case, the agent-based equivalent model will enable us to move beyond the MaxEnt analysis in the sense of easily expanding the model to take into account heterogeneous agents.

Similar to the previous agent-based example in Section 3, let us study a system composed by $N = 1000$ firms, each of them characterized by a binary variable setting their investment decision in a specific market or sector, $a_i = \{0, 1\}$, and let $\bar{a}_t = \sum_{i=1}^N a_i / N$ be the average action or fraction of

firms investing at a given point in time t . A quantal action with a market feedback mechanism can be easily modeled in this agent-based setting; however, different models can be defined based on different specifications of the market feedback function. As an example, we will explore here a feedback function in which the magnitude of the correction on the rate of profit x_t is proportional to the number of new firms entering the market,

$$x_{t+1} = x_t - \alpha (\bar{a}_t - \bar{a}_{t-1}) N, \quad (20)$$

with α being a parameter setting the strength of the market feedback mechanism. As in the previous agent-based example in Section 3, it is useful to specify the time evolution of the system as a set of actions and updates to individual and aggregate state variables taking place at every time step. In particular, let us start by assuming an initial rate of profit x_0 . Using the same decision rule for individual firms as in the MaxEnt example above,

$$P(a_i = 1|x_t) = \frac{1}{1 + e^{\frac{\mu - x_t}{T}}}, \quad (21)$$

where μ is, as before, the hurdle rate of profit, and using the initial value of the rate of profit x_0 , each of the N firms can find its initial probability of investing. Then a random number is uniformly drawn between 0 and 1 for each firm, who will decide to invest if this number is below the corresponding probability. Once the new individual states are computed, an updated value of the rate of profit can be found using equation (20). Finally, the following process can be repeated for a number of time steps: a firm is chosen at random, its new probability of investing is computed, a new random number is drawn, the firm's market position is updated accordingly, and a new profit rate is found. Note that, since only one firm is updated per time step, the described process constitutes an asynchronous update mechanism, chosen here in order to avoid unreasonable profit rate fluctuations and overshootings. Figure 7 shows the resulting distribution of the profit rate with different behavior temperatures T and different market feedback strengths α .




Figure 7: Distribution of the rate of profit for a fixed market feedback strength α and different behavior temperatures T (on the left) and for fixed behavior temperature T and different values of α (on the right).

The left-hand side plot in Figure 7 shows different distributions of the profit rate with varying behavior temperature T and constant market feedback strength $\alpha = 0.001$, while the right-hand side plot shows the distributions with varying α and constant $T = 1.00$. The agent-based model replicates the impact of the temperature T on the profit distribution that we obtained in the previous subsection with the MaxEnt model: the higher the behavioral temperature T the more dispersed the distribution. Curiously enough, the market feedback effect has quite the opposite impact in the agent-based model. In the MaxEnt, the higher the market feedback (low ξ and high β), the more peaked the distribution. In the agent-based model, higher market feedback (high α) leads to a more dispersed distribution. The main reason behind this difference is that, while the implicit market feedback effect ξ in the MaxEnt determines the degree of convergence of the profit rate to the central tendency, the feedback effect α in the agent-based model determines the size of the market response to the action. Therefore, a higher feedback effect in the agent-based model can lead to overshooting, and thus to more volatile fluctuations around the mean and a higher dispersion of the resulting distribution.

As explained above, one of the most significant advantages of agent-based models is their flexibility and the ease with which they can be extended to take into account different kinds of heterogeneity. As an illustration, we have extended the previously described model so as to include two types of agents with different behavior temperatures. In particular, we consider a system composed of $N/2$ agents of type 1, with behavior temperature $T_1 = 0.01$, and $N/2$ agents of type 2, with behavior temperature $T_2 = 10$. In order to understand the effects of this heterogeneity, it is useful to define the average profit rate of agents of type k at time step t as $\bar{x}_{t,k} = \bar{a}_{t,k} \cdot x_t$, where $\bar{a}_{t,k}$ is simply the average action of agents of type k at time t . Figure 8 shows the resulting distribution of the average profit rate for each type of agent, $P(\bar{x}_{t,1})$ and $P(\bar{x}_{t,2})$, along with the distribution of the profit rate for the whole system of firms, $P(x_t)$.




Figure 8: Distributions of different measures of the rate of profit for a fixed market feedback strength $\alpha = 0.01$, two types of agents with behavior temperatures $T_1 = 0.01$ and $T_2 = 10$, and for an initial profit rate below the firms' hurdle rate ($x_0 < \mu$, on the left) and above the firms' hurdle rate ($x_0 > \mu$, on the right). The dotted line represents the distribution of the aggregate market profit rate, while the solid and the dashed lines represent the distributions for the average profit rate for the low (T_1) and high-temperature (T_2) agents respectively.

Given that both types of agents have exactly the same probability of investing when the market rate coincides with the firms' hurdle rate, $x_t = \mu$, and very similar probabilities around that point, it is more interesting to study the case where both rates are significantly different. To this end, we initialize the system with two different values for the initial rate of profit⁷, one below the firms' hurdle rate (left-hand side panel in Figure 8) and the other one above it (right-hand side panel in Figure 8). When market rates are below the firms' hurdle rate, as is the case in the left-hand side panel, the lower temperature agents rarely invest and therefore they have a very peaked distribution around zero average profit rate. In contrast, the higher temperature agents randomly invest even when the profit rate is below their hurdle rate and therefore their average profit rate distribution is centered around a negative value. When the market profit rate is persistently above the firm's hurdle rate, the low-temperature agents almost always invest and exploit the higher profit rate more often than the higher temperature agents, who invest more randomly. Consequently, the profit rate distribution of the low-temperature agents has a higher mode. By combining the insights from both plots, we can say that low-temperature agents make less losses than high-temperature agents when the profit rate is significantly below the hurdle rate (as in the left-hand side plot), while they make higher profits when the rate is significantly above the hurdle rate (as in the right-hand side plot). This is consistent with the fact that high-temperature agents are subject to less informational uncertainty than low-temperature agents, and thus can make better informed decisions.

The MaxEnt model of market feedback with heterogeneous agents has not been explored yet partly due to mathematical difficulty. In this regard, the results of the agent-based equivalent model presented here can provide important insights for a potential MaxEnt model extension. First, it is noticeable that

⁷Note that, due to the definition of the market feedback mechanism in equation (20), the market rate fluctuates around its initial value.

the unimodal pattern of the aggregate distribution stays the same even when two different types of agents operate in the market and have a completely different realization of their own profit rate. Heterogeneous levels of behavior temperature determine how often each type of agent enters the market but doesn't change the underlying market feedback mechanism that predominantly determines the unimodal shape of the distribution. Second, even when a given type of agent has a high behavior temperature, its own realization of the profit rate can have a narrow dispersion. The second plot in Figure 8 shows that the high temperature agents have a narrower distribution of the realized profit rate when the market has a persistently positive profit rate. This is due to the fact that the high-temperature agents do not always respond to a very high expected profit rate and thus fail to form enough mass on the right tail.

5 Conclusions

This paper examines MaxEnt and agent-based modeling with the aim to show how these two different analytical tools can complement each other and can provide a powerful conceptual/empirical tool for the analysis of complex economic problems. We argued that the MaxEnt approach is particularly useful in understanding the qualitative properties of systems in quasi-equilibrium due to its closed-form solutions, which could serve as guidance to developing first-approximation agent-based models. In contrast, agent-based modeling is highly flexible and can easily accommodate a wide choice of behavioral and institutional assumptions due to its simulation-based approach. In this regard, an agent-based model can be implemented based on a MaxEnt model and provide a richer picture of the latter model, potentially guiding further extensions.

We have demonstrated the complementarity of MaxEnt and agent-based modeling using two simple economic models: a quantal response social interaction model and a market feedback model. In illustrating each of them, we have first derived the statistical equilibrium using the MaxEnt approach, after which an agent-based version of the model has been implemented. In particular, we have highlighted the capabilities of agent-based modeling for tracking irreversible and out-of-equilibrium dynamic processes, using the social interaction case as an example, as well as for exploring the consequences of agent heterogeneity, using the market feedback model as an example.

References

- Anas, A. (1983). Discrete choice theory, information theory and the multinomial logit and gravity models. *Transportation Research Part B: Methodological*, 17(1):13 – 23.
- Axtell, R. (2000). Why agents? on the varied motivations for agent computing in the social sciences. *The Brookings Institution, Center on Social and Economic Dynamics Working Paper, Nov, 17*.
- Bais, F. A. and Farmer, J. D. (2007). *The Physics of Information*.
- Baptista, R., Farmer, J. D., Hinterschweiger, M., Low, K., Tang, D., and Uluc, A. (2016). Macroprudential policy in an agent-based model of the uk housing market.
- Blackwell, K. (2018). Entropy constrained behavior in financial markets: A quantal response statistical equilibrium approach to financial modeling.
- Borwein, J. M. and Lewis, A. S. (1991). On the convergence of moment problems. *Transactions of the American Mathematical Society*, 325(1):249–271.
- Boyd, S. and Vandenberghe, L. (2004). *Convex Optimization*. Cambridge University Press.
- Castellana, M. and Bialek, W. (2014). Inverse spin glass and related maximum entropy problems. *Phys. Rev. Lett.*, 113:117204.
- Cherkashin, D., Farmer, J. D., and Lloyd, S. (2009). The reality game. *Journal of Economic Dynamics and Control*, 33(5):1091 – 1105. Complexity in Economics and Finance.
- Cover, T. M. and Thomas, J. A. (2006). *Elements of information theory (2nd Edition)*. Wiley-Interscience.
- Foley, D. K. (2017). Information theory and behavior. Working Papers 1731, New School for Social Research, Department of Economics.
- Irle, A., Milakovic, M., Alfarano, S., and Kauschke, J. (2008). A statistical equilibrium model of competitive firms. Economics Working Papers 2008-10, Christian-Albrechts-University of Kiel, Department of Economics.
- Jaynes, E. T. (2003). *Probability Theory: The Logic of Science*. Cambridge University Press.
- MacKay, D. J. (2005). *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press.
- McKelvey, R. D. and Palfrey, T. R. (1996). Quantal response equilibria for normal form games.
- McKelvey, R. D. and Palfrey, T. R. (1998). Quantal response equilibria for extensive form games. *Experimental Economics*, 1(1):9–41.

- Sakellariou, J., Tria, F., Loreto, V., and Pachet, F. (2017). Maximum entropy models capture melodic styles. *Scientific Reports*, 7(1):9172.
- Scharfenaker, E. and dos Santos, P. L. (2015). The distribution and regulation of tobin's q. *Economics Letters*, 137:191–194.
- Scharfenaker, E. and Foley, D. K. (2017). Quantal response statistical equilibrium in economic interactions: Theory and estimation. *Entropy*, 19(9).
- Scharfenaker, E. and Semieniuk, G. (2017). A statistical equilibrium approach to the distribution of profit rates. *Metroeconomica*, 68(3):465–499.
- Shannon, C. (1948). A mathematical theory of communication. *Bell System Technical Journal*, 27:379–423.
- Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, 50(3).
- Sims, C. A. (2006). Rational inattention: A research agenda. Available at <http://sims.princeton.edu/yftp/RIplus/RatInattPlus.pdf>.
- von Neumann and Morgenstern (1944). *Theory of Games and Economic Behavior*. Princeton University Press.
- Xu, Z., Bi, G., Zhou, D., and Cai, D. (2017). A dynamical state underlying the second order maximum entropy principle in neuronal networks. *Communications in Mathematical Sciences*, 15(3):665–692.
- Yang, J. (2018). Information theoretic approaches in economics. *Journal of Economic Surveys*, 32(3):940–960.