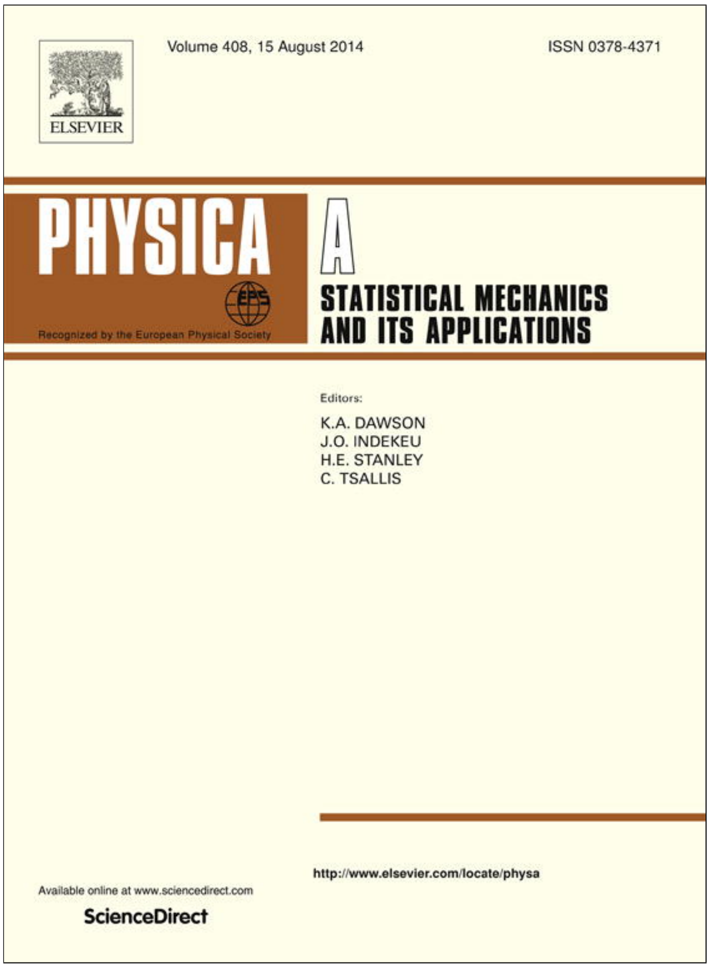
Volume 408, 15 August 2014
ISSN 0378-4371
ELSEVIER
PHYSICA
Recognized by the European Physical Society
A
STATISTICAL MECHANICS
AND ITS APPLICATIONS
Editors:
K.A. DAWSON
J.O. INDEKEU
H.E. STANLEY
C. TSALLIS
http://www.elsevier.com/locate/physa
Available online at www.sciencedirect.com
ScienceDirect

Physica A 408 (2014) 149–163

Contents lists available at ScienceDirect

Physica A

journal homepage: www.elsevier.com/locate/physa

An entropic framework for modeling economies

Ariel Caticha [a], Amos Golan [b,*]

[a] Department of Physics, University at Albany-SUNY, Albany, NY 12222, USA

[b] Department of Economics, Info-Metrics Institute, American University, Washington, DC, USA

HIGHLIGHTS

- An information-theoretic framework for modeling an economic system is developed.
- Statistical equilibrium is derived from fundamental principles via the method of maximum entropy.
- Minimal information, at the macroscopic level, is introduced.
- Prices arise naturally as statistical constructs via the Lagrange multipliers.
- That framework can be used for modeling other social systems and networks.

ARTICLE INFO

Article history:
Received 12 September 2013
Received in revised form 3 April 2014
Available online 13 April 2014

Keywords:
Economic statistical equilibrium
Information
Entropy
Modeling

ABSTRACT

We develop an information-theoretic framework for economic modeling. This framework is based on principles of entropic inference that are designed for reasoning on the basis of incomplete information. We take the point of view of an external observer who has access to limited information about broad macroscopic economic features. We view this framework as complementary to more traditional methods. The economy is modeled as a collection of agents about whom we make no assumptions of rationality (in the sense of maximizing utility or profit). States of statistical equilibrium are introduced as those macrostates that maximize entropy subject to the relevant information codified into constraints. The basic assumption is that this information refers to supply and demand and is expressed in the form of the expected values of certain quantities (such as inputs, resources, goods, production functions, utility functions and budgets). The notion of economic entropy is introduced. It provides a measure of the uniformity of the distribution of goods and resources. It captures both the welfare state of the economy as well as the characteristics of the market (say, monopolistic, concentrated or competitive). Prices, which turn out to be the Lagrange multipliers, are endogenously generated by the economy. Further studies include the equilibrium between two economies and the conditions for stability. As an example, the case of the nonlinear economy that arises from linear production and utility functions is treated in some detail.

1. Introduction: an entropic framework

We wish to model an economy that consists of agents – individuals or firms – and goods that are produced by some agents and consumed by others. The agents may differ from each other in what they produce, in how productive they are,

* Corresponding author. Tel.: +1 2028853783.
E-mail addresses: ariel@albany.edu (A. Caticha), agolan@american.edu (A. Golan).

http://dx.doi.org/10.1016/j.physa.2014.04.016
0378-4371/

and also in their different preferences as to what goods they consume. In the more traditional approaches to economic modeling, one takes the point of view that the agents are engaged in a game of interactions. Each agent strives to maximize its utility function. It has full knowledge of its own utility function and some possibly partial knowledge about the utilities and preferences of the other agents. Under this modeling approach one is able to infer the broader macroscopic features of the economy by keeping track of the detailed behavior of its individual agents.

In this work, we take a complementary approach. We study the economy as if from a bird's eye point of view. We imagine an external observer – we are the observer – whose goal is to model the economy in order to make inferences. We collect information that is accessible at the macroscopic level but we have a very limited access to information about the detailed actions of the various agents. For example, we will assume that the agents have preferences but no assumptions will be made about whether they behave rationally or not—that is, we do not assume that the agents strive to maximize utility and profit at each moment of time.

We propose an information-theoretic framework for modeling economies formulated as a problem of inference with limited information. The basic idea – the method of Maximum Entropy – is well established [1–4]. The economic states of statistical equilibrium are those that maximize entropy subject to relevant information that is introduced through constraints. Prices are not introduced directly, but are rather determined via the Lagrange multipliers that indirectly reflect the information on supply and demand. There are many previous attempts to apply entropic methods to problems in economics. They range, to name a representative sample where further references can be found, from data analysis and econometrics [3,5], thermodynamic analogies [6,7], to modeling the statistical equilibrium of an ensemble of producing firms [8], or of an ensemble of transactions or offer sets [9–11], and the econophysics of income distributions [12,13].

Although our objective is to develop a framework for modeling and not to establish analogies with other systems, it is interesting to note that such analogies arise naturally. The obvious analogy is with thermodynamics and statistical mechanics. (see, *e.g.*, Ref. [7].) We hasten to point out that the analogies arise not because there is any underlying similarity between physical and economic systems—there is none. It is rather that in both cases we reason with extremely incomplete information—too many atoms following unknown trajectories in one case and too many agents making unknown decisions in the other. It is the entropic method of inference that leads to analogies by imposing rigid constraints on the formal mathematical structure. In both cases we deal with constrained maximization and the analysis involves the introduction of the Lagrange multipliers and the Legendre transforms.

The notion of statistical equilibrium that we adopt here differs in important respects from other notions of equilibrium proposed elsewhere (for earlier work and references see Refs. [9,8,10]). Here it is the fundamental entities, namely the agents, in terms of which the equilibrium is defined. This allows us to simultaneously discuss exchange and production, including the effects of competitiveness and the distribution of goods, within the same economic model.

We expect that the framework presented here will be pragmatically useful. The reason is that, just as in statistical mechanics, from the point of view of someone who wants to understand the broad macroscopic features of the system most of the "microscopic" details will turn out to be irrelevant because they will be washed out when taking averages. Any calculational method that manages to ignore those irrelevant microscopic details from the beginning will therefore represent a huge gain in efficiency. Thus, our framework seeks to focus directly on the information that matters while systematically disregarding all those details that do not. The fundamental challenge lies, of course, in identifying correctly those specific pieces of information that turn out to be relevant.

In Section 2 the model is introduced with the definitions and discussion of the quantities of interest (agents, goods, production, utilities and budgets). The relevant information on the basis of which all inferences will be carried out – information about consumption and production of various goods, and about the preferences and budgets of agents – is discussed in Section 3. In Section 4 we introduce the central concept of statistical economic equilibrium and use the method of maximum entropy to find the probability distribution of the goods allocated to each agent. The meaning of the statistical equilibrium is further discussed in Section 5 where we introduce the concept of economic entropy. Then, in Section 6, we establish the relation between the Lagrange multipliers resulting from entropy maximization and the economic prices. In Section 7 we discuss the conditions for equilibrium between two economies and derive the conditions for such an equilibrium to be stable. In Section 8, as an illustration of the method, we discuss the special case of an economy characterized by linear production and utility functions. We find that despite linearity at the level of the individual agent, the resulting economy is highly non-trivial. We also discuss the conditions under which the economy moves toward a competitive market or toward a more concentrated market. Finally, in Section 9 we conclude with a summary and discussion of future research.

2. The fundamental entities

Any problem of inference rests on several types of assumptions. First we must decide on the subject matter. This involves assumptions about the fundamental entities that define the system including those quantities we want to infer.

Second, as discussed in the next section, one must make assumptions about the relationships among the fundamental entities—the quantities that are observable, those that are to be inferred, and the relevant information that links them.

Our system is an economy that consists of individuals and firms, called agents, and goods that are produced by some agents and consumed by others. The agents will be labeled $a = 1 \ldots A$ and the goods will be labeled $g = 1 \ldots G$.

The agents may differ from each other in a variety of ways. We will model the effects of two types of differences.

First, differences in human capital and opportunities lead agents to adopt different technologies or different business models and therefore they differ in what they produce and in how productive they are. Second, the various agents have different preferences as to what goods they consume.

Some of the consumed goods are used as input resources (*e.g.*, raw materials, intermediate goods, energy, and services) required for production while others are required for the daily business of living (*e.g.*, food, clothes, and books). For any agent a let f_{ag} be the amount of good g produced by a, let x_{ag} be the amount of a good g used as input for production, and let y_{ag} the amount used in general household consumption.

Let $x_{agg'}$ be the amount of resources (inputs) of type g used by agent a for the production of good g'. The amount of g' produced by a is, naturally, a function of the inputs/resources consumed by a,[1]

$$f_{ag'} = f_{ag'}(x_{a1g'}, x_{a2g'} \dots x_{aGg'}). \tag{1}$$

Therefore, the total amount of good g used for production by agent a is

$$x_{ag} = \sum_{g'} x_{agg'}. \tag{2}$$

The scheme is flexible enough to describe situations where not all agents produce the same goods or require the same resources. For example we expect that for any given agent a most of the production goods $f_{ag'}$ will vanish identically—which means that agent a does not produce the particular good g'. The set of quantities $\{x_{agg'}\}$ can be quite large – it has $A \times G^2$ elements – but we expect that most of them will vanish. Indeed, whenever a particular input g is not needed for the production of a good g' the production function $f_{ag'}$ is independent of the corresponding argument $x_{agg'}$ and one can safely expect that agent a will not need to purchase it. Thus, we can set

$$x_{agg'} = 0 \quad \text{whenever } \frac{\partial f_{ag'}}{\partial x_{agg'}} = 0. \tag{3}$$

For example, the good g' might be money (or an initial endowment) and the agent might "produce" it by merely withdrawing it from a bank (*e.g.*, the agent lives off its endowment). This is an instance where the production $f_{ag'}$ requires no input ($x_{agg'} = 0$). Another possibility is that a fixed amount of a certain good $x_{agg'}$ – say a chemical catalyst – is sufficient for production of any amount of g'. All such special cases will require special attention.

Goods must be bought and sold. This brings us to the second way that agents may differ from each other, namely, through their different preferences for various consumption goods. The relative preferences of each agent for a bundle of goods $\{y_{ag}\}$ is described by the utility function,

$$u_a = u_a(y_{a1}, y_{a2} \dots y_{aG}). \tag{4}$$

In the conventional microeconomic analysis it is assumed that each rational agent will choose its bundle of goods $\{y_{ag}\}$ so as to maximize the utility subject to budget constraints. In our entropic framework utility functions are introduced because they convey relevant information but we make no assumption about agents seeking to maximize their individual utilities.

Several pieces of information are required in order to specify an economy. Part of this information is contained in the production functions f_{ag} and in the individual preference functions u_a. We will assume these functions are given—they specify the model.

A state of the economy where the distribution of goods among the various agents is described at the micro level (in full microscopic detail) is called a *microstate*. In the present scheme a microstate is defined by a detailed specification of the distribution of all inputs and goods,

$$\{x_{agg'}\} = x \quad \text{and} \quad \{y_{ag}\} = y. \tag{5}$$

In contrast, a *macrostate* of the economy is a state for which we only have partial information and which, therefore, is described by a probability distribution

$$P(\{x_{agg'}, y_{ag}\}) = P(x, y). \tag{6}$$

Our objective is to assign $P(x, y)$ on the basis of the available information which is very limited.

This framework of agents and goods that are exchanged in unspecified ways is sufficiently broad to tackle a wide variety of situations of equilibrium. In particular, money can be included as a very special good which figures prominently in the utility functions of all agents, and which is special in that it can be produced or consumed by some special agents that we call Central Banks. The microstate of any agents is characterized by its complete stock of goods, which certainly includes money, gold, bonds, art, and anything else that has value.[2]

1 We adopt the usual convention of denoting the amount produced and the production function by the same symbol, $f_{ag'}$.

2 Money is clearly an important topic but in this first paper that describes the general framework we did not want to go into the technical details of banking and all the various forms that money can take in a modern economy—currency, checking accounts, money market funds, gold, bonds, credit cards, and all sorts of other financial instruments.

3. The relevant information

Next we discuss the relevant information that defines the macrostates.

3.1. Consumption

The first piece of macro-information is that the total amount of a good g whether used as input for production or for household consumption,

$$\sum_{ag'} x_{agg'} = \sum_{a} x_{ag} \quad \text{and} \quad \sum_{a} y_{ag}, \tag{7}$$

is always finite and limited by nature. Instead of sharp constraints on these quantities it turns out to be convenient to impose constraints on their expected values. Let

$$\left\langle \sum_{ag'} x_{agg'} \right\rangle = \left\langle \sum_{a} x_{ag} \right\rangle = X_g \quad \text{and} \quad \left\langle \sum_{a} y_{ag} \right\rangle = Y_g. \tag{8}$$

The quantities X_g and Y_g are, in principle, observable. In what follows the variables that define the microstate of the economy (x, f, y) will be denoted with lower case latin letters; their expected values which define the macrostate will be in upper case (X, F, Y); and the corresponding Lagrange multipliers will be Greek letters, (α, β, γ).

3.2. Production

The second piece of macro-information is that the total production of each good g, at a given time, is constrained either by natural causes or by the economic environment or by both. We impose a constraint on the expected total production, Eq. (1),

$$\left\langle \sum_{a} f_{ag} \right\rangle = F_g. \tag{9}$$

The quantities F_g are, in principle, observable.

3.3. Supply equals demand: production equals consumption

In a closed economy the constraints (8) and (9) are not independent. All consumed goods must have been produced and conversely all produced goods must be consumed. Therefore, the expected aggregate amounts of each good g must satisfy

$$F_g = X_g + Y_g. \tag{10}$$

The production equals consumption constraint (10) is sufficiently flexible to cover a wide variety of special cases. It may happen that it is not natural to conceive of a particular good g as a input for production but only as a good for household consumption. For such goods we can set $X_g = 0$ and the equation above reads $F_g = Y_g$. On the other hand there may exist other goods that are naturally conceived as inputs for production and not for household consumption. In these cases we set $Y_g = 0$ so that $F_g = X_g$.

3.4. Preferences and budgets

Next we account for two additional factors. First, we want to represent the uneven preferences of different agents for different goods. Information of this kind is captured by a constraint on the expected utility for each agent,

$$\langle u_a \rangle = U_a. \tag{11}$$

The quantity U_a provides some measure of how satisfied an agent is in a particular economy. In order to compare the utilities of one agent with another it would be necessary to adopt some normalization criteria. Such criteria are not unique and different choices may turn out to be convenient for different purposes.

There is no implication that individual agents will necessarily behave "rationally" in the sense of attempting to maximize $\langle u_a \rangle$; they might or might not. The models discussed here are silent about the issue of rational behavior—which might, in the end, be an important and welcome feature. See also [14].

The quantities X_g, F_g, and Y_g represent expected total amounts of the good g and are in principle directly observable. (Strictly speaking it is the total amounts that are directly observable and not their expected values. For sufficiently large systems, however, the probability that these two quantities differ becomes negligible.) The situation with the quantities U_a is subtler: we do not actually know their numerical values. But then how can we justify invoking an unknown U_a as a piece of relevant information? The answer is that what justifies constraints on $\langle u_a \rangle$ is not that their values happen to be known, but rather that they represent relevant information that *ought* to be known and that must be included in the analysis in order to make successful predictions. The procedure is to impose such constraints as if the values U_a were known and the result is a family of macrostates $P(x, y|U_a)$ where U_a are parameters that remain to be determined.[3] The additional information required to pin down the U_a is conveyed through the budget constraints discussed next.

We also need to capture the fact that different agents have access to different budgets. It is not sufficient to impose that for every good the overall demand must match the overall supply, Eq. (10); we must also impose that for each individual agent its capacity to consume must to some extent be constrained by its ability to produce (its budget) or its overall endowment. Thus, the expected utility $\langle u_a \rangle = U_a$ of an agent must be in some way correlated with how productive (and therefore wealthy) the agent is. For each agent a we impose that the expected total cost of consumption must match the value of production,

$$\sum_g \tilde{\pi}_g \langle f_{ag} - x_{ag} \rangle = \sum_g \tilde{\pi}_g \langle y_{ag} \rangle, \tag{12}$$

where $\tilde{\pi}_g$ is the price of good g.[4,5]

It is important to note that the budget match between production and consumption need not be perfect but holds only on the average, in expected value, so that there is room for deviations (fluctuations) about the mean. We are only interested in the broad picture that models the economy as a whole and we do not assume that individual agents are rigorously constrained by their budgets.

This concludes the specification of the relevant information used here. We note, in particular, that no additional information is needed to determine the prices. We will show that the prices $\tilde{\pi}_g$ will turn out to be simply related to the Lagrange multipliers associated with Eqs. (8) and (9).[6] Their relation to the expected utilities U_a is one of self-consistency: given the utilities U_a the prices $\tilde{\pi}_g$ are determined by the constraints on supply and demand. On the other hand, given the prices $\tilde{\pi}_g$, the utilities U_a are determined by budget constraints Eq. (12).

4. Statistical economic equilibrium

In this section the macrostate $P(x, y)$ will be determined using the method of maximum entropy. The result is that the distribution $P(x, y)$ will incorporate the information codified into the chosen constraints while remaining maximally non-informative about everything else.

4.1. The basic assumption

Concerning the distribution $P(x, y)$ one could adopt the conservative or informationally weaker point of view that $P(x, y)$ has merely predictive value: it allows one to make the best inferences that are justified by the limited information at hand. Unfortunately, however, the best inference need not be any good. There are no guarantees; it all depends on the quality of the information.

Instead, we will adopt the more optimistic or informationally stronger point of view that the set of relevant constraints that fully characterizes the system under study has already been identified. This assumption will be adopted throughout the rest of this paper.

Completeness Assumption: The chosen constraints on expected production, consumption, utilities, and budgets form a complete set in the sense that they capture all the information that is relevant for the prediction of certain macroscopic features of economic interest.

The question then becomes: under what conditions can we reasonably expect this assumption to hold? Suppose that we found that the quantities X, F, Y and U_a depend on time. We would immediately conclude that additional information about their temporal behavior – perhaps in a form given by ''equations of motion''– would be highly relevant and should have been included in the analysis. Therefore the completeness assumption fails in time-dependent situations. Conversely, when the completeness assumption holds the situation must be statistical equilibrium.

3 An analogous problem occurs in statistical mechanics where thermodynamic equilibrium requires a constraint on expected energy $\langle E \rangle$. The value of $\langle E \rangle$ is not known a priori but such a constraint is used anyway to determine a family of distributions – the Boltzmann distributions – in which temperature appears as a parameter. The temperature is then either directly measured with a thermometer or indirectly inferred from other data.

4 The budget constraint applies at the microscopic level. This is somewhat analogous to energy conservation in statistical mechanics. It is not enough to require that there be an overall energy conservation law; it is necessary that the conservation apply to every microscopic interaction between molecules. Here, however, the budget constraint is weaker; it holds only in expectation.

5 A common convention is to use different symbols to distinguish input and output prices. Such distinctions are relative to any given agent. Our notation is better adapted to a description of the economy as a whole where any given good can be an input for some agents and an output for others.

6 There is a useful analogy to physics here. The concepts of temperature and chemical potential cannot be imported from mechanics into thermodynamics. They are emergent concepts of purely statistical origin. Just as temperature regulates the flow of heat, and a chemical potential regulates the flow of particles, we shall later see that prices regulate the flow of goods.

The macrostate described by the probability distribution obtained on the basis of the above information – and nothing else – is defined as the state of *economic statistical equilibrium*. In our framework this macrostate is the probability distribution that maximizes entropy.

The notion of equilibrium that we have in mind represents a useful idealization but is perhaps strictly unattainable in practice. The notion of statistical equilibrium might be approximately valid over time scales in an intermediate range. On one hand the time scales must be sufficiently long to allow the system to settle down into some approximate equilibrium. On the other hand if we wait too long the statistical equilibrium will tend to drift due to the appearance of new technologies, of new sources of raw materials, of shifts in preferences, fashions, and social structures.

Finally, it is important to draw a sharp distinction between the macroscopic statistical equilibrium that is proposed here and is derived by maximizing entropy from the other more traditional notions of equilibrium such as Walrasian or Nash equilibria that operate either at the microeconomic level of individual agents or at the collective behavior of agents engaged in trades. Our approach complements these more traditional approaches and is based on a smaller set of assumptions/information.

4.2. The economic macrostate

The macrostate distribution $P(x, y)$ is obtained by maximizing the entropy subject to the constraints (8)–(11) plus normalization. The entropy $\mathcal{S}[P, Q]$ of the distribution $P(x, y)$ relative to a reference or "prior" distribution $Q(x, y)$ is

$$\mathcal{S}[P, Q] = -\sum_{x,y} P(x, y) \log \frac{P(x, y)}{Q(x, y)}. \tag{13}$$

First we discuss the assignment of the "prior" distribution $Q(x, y)$. Earlier we argued that agents differ either through the production functions f_{ag}, or the preferences u_a and that this information is entered through the constraints. We design our model so that a priori – that is, in the absence of the information introduced through the constraints – the agents will be identical in all respects. Since the result of an unconstrained maximization of $\mathcal{S}[P, Q]$ is that the posterior coincides with the prior, $P = Q$, we require that the prior $Q(x, y)$ represents a uniform distribution of goods and resources.

The choice of a uniform Q, which represents a welfare state and a fully competitive market, is significant in that it affects the economic interpretation of the entropy: macrostates $P(x, y)$ with higher entropy $\mathcal{S}$ reflect a more uniform distribution of resources and goods. This is an important topic to which we will return later. The generalization to a non-uniform Q is straightforward but will not be discussed here.

Maximizing $\mathcal{S}[P, Q]$ for uniform Q is equivalent to maximizing the Gibbs–Shannon entropy

$$\mathcal{S}[P] = -\sum_{x,y} P(x, y) \log P(x, y). \tag{14}$$

The result of maximizing $\mathcal{S}[P]$ subject to the five constraints (8)–(11) plus normalization is

$$P(x, y) = \frac{1}{\zeta} \exp\left[-\sum_{ag} (\alpha_g x_{ag} - \beta_g f_{ag} + \gamma_g y_{ag}) - \sum_a \mu_a u_a \right] \tag{15}$$

where α_g, β_g, γ_g, and μ_a are the Lagrange multipliers, ζ is the normalization constant or partition function, and x_{ag}, f_{ag}, and u_a are given by Eqs. (2), (1), and (4). The signs of the multipliers are chosen for later convenience. The partition function ζ is a function of the Lagrange multipliers

$$\zeta(\alpha, \beta, \gamma, \mu) = \sum_{xy} \exp\left[-\sum_{ag} (\alpha_g x_{ag} - \beta_g f_{ag} + \gamma_g y_{ag}) - \sum_a \mu_a u_a \right]. \tag{16}$$

The Lagrange multipliers are determined by

$$-\frac{\partial \log \zeta}{\partial \alpha_g} = X_g, \qquad \frac{\partial \log \zeta}{\partial \beta_g} = X_g + Y_g, \tag{17}$$

$$-\frac{\partial \log \zeta}{\partial \gamma_g} = Y_g, \qquad -\frac{\partial \log \zeta}{\partial \mu_a} = U_a, \tag{18}$$

which are conveniently summarized as

$$\delta \log \zeta = \sum_g \left[X_g \delta(\alpha_g - \beta_g) + Y_g \delta(\gamma_g - \beta_g) \right] - \sum_a U_a \delta \mu_a. \tag{19}$$

An alternative choice of variables that will turn out to be convenient is the following. Eliminate X_g and Y_g in favor of their sum and difference

$$F_g = Y_g + X_g \quad \text{and} \quad V_g = Y_g - X_g, \tag{20}$$

or,

$$X_g = \frac{1}{2}(F_g - V_g) \quad \text{and} \quad Y_g = \frac{1}{2}(F_g + V_g). \tag{21}$$

Then (19) can be written as

$$\delta \log \zeta = -\sum_g (F_g \delta\pi_g + V_g \delta\lambda_g) - \sum_a U_a \delta\mu_a, \tag{22}$$

where we introduced new multipliers,

$$\pi_g = \frac{1}{2}(\alpha_g + \gamma_g) - \beta_g \quad \text{and} \quad \lambda_g = \frac{1}{2}(\gamma_g - \alpha_g). \tag{23}$$

The introduction of the new variable V_g and its associated multiplier λ_g are dictated by the mathematical formalism. Their economic significance will be discussed later. The ratio

$$\frac{V_g}{F_g} = \frac{Y_g - X_g}{Y_g + X_g} \tag{24}$$

is easy to interpret: it ranges from the extreme value -1, which means that the good g is used purely as an input for production, to the opposite extreme $+1$, which means that the good g is used purely for consumption. Just as the multiplier π_g regulates the overall amount of F_g the multiplier λ_g regulates the value of V_g which captures the input to consumption mix.

4.3. The distributions of goods for individual agents

The macrostate $P(x, y)$ takes the form of a product over the x and the y variables and also over different agents,

$$P(x, y|\alpha, \beta, \gamma, \mu) = \prod_a P_a(x_a|\alpha, \beta) P_a(y_a|\gamma, \mu_a) \tag{25}$$

where x_a and y_a are shorthand for the sets $\{x_{agg'}\}_a = x_a$ and $\{y_{ag}\}_a = y_a$ for a fixed agent a, and

$$P_a(x_a|\alpha, \beta) = \frac{1}{\zeta_a^x} \exp\left[-\sum_g (\alpha_g x_{ag} - \beta_g f_{ag})\right] \tag{26}$$

$$P_a(y_a|\gamma, \mu_a) = \frac{1}{\zeta_a^y} \exp\left[-\sum_g \gamma_g y_{ag} - \mu_a u_a\right], \tag{27}$$

and $\zeta_a^x(\alpha, \beta)$ and $\zeta_a^y(\gamma, \mu_a)$ are the appropriate normalization constants.

Thus, conditional on the Lagrange multipliers the inputs x_a and the household goods y_a consumed by different agents behave like uncorrelated independent systems. This means that in this model individual agents affect each other only indirectly through their interaction with the economy as a whole. This phenomenon is consistent with the more traditional economic modeling.[7]

The various expected values for agent a are

$$-\frac{\partial \log \zeta_a^x}{\partial \alpha_g} = \langle x_{ag}\rangle = X_{ag}, \qquad \frac{\partial \log \zeta_a^x}{\partial \beta_g} = \langle f_{ag}\rangle = F_{ag}, \tag{28}$$

$$-\frac{\partial \log \zeta_a^y}{\partial \gamma_g} = \langle y_{ag}\rangle = Y_{ag}, \qquad -\frac{\partial \log \zeta_a^y}{\partial \mu_a} = U_a. \tag{29}$$

which are summarized as

$$\delta \log \zeta_a^x = \sum_g (-X_{ag}\delta\alpha_g + F_{ag}\delta\beta_g), \tag{30}$$

$$\delta \log \zeta_a^y = -\sum_g Y_{ag}\delta\gamma_g - U_a \delta\mu_a. \tag{31}$$

The economic interpretation of the μ_a can be read from the distribution $P_a(y_a|\gamma, \mu_a)$. Eq. (27) shows that the probability of microstates with higher utility is exponentially suppressed but not at the same rate for all agents. A lower value of μ_a implies a higher probability of attaining microstates of higher utility. Conversely, agents with higher μ_a are effectively confined

7 One could extend the model to include interactions among agents leading to non-trivial correlations by imposing additional constraints. It is an important (and positive) feature of the entropic inference framework that the model will not include such interactions unless they are deliberately introduced into the model.

to microstates with low utility. It is thus clear that agent *a* will prefer low values of μ_a over high values. Indeed, as we discuss in Section 6.1 and then explicitly show in Section 8, an agent that manages to increase its productivity (for example, by improving its human capital) will, other things being equal, achieve a lower value of μ_a and will become effectively wealthier.

Deciding whether an agent *a* with a low μ_a is wealthier in some absolute sense than some other agent a' with a higher $\mu_{a'}$ depends on comparing the utility functions of one agent with another. Such absolute comparisons are not possible because utility functions measure relative preferences for bundles of goods; they are determined only up to an arbitrary agent-dependent scale factor k_a. Fortunately such arbitrary scale factors have no observable effects because, as we see in Eq. (27), the distribution $P_a(y_a|\gamma, \mu_a)$ depends on μ_a and on u_a only through the product $\mu_a u_a$, which implies that if utilities are scaled by a factor k_a the corresponding multipliers are scaled by k_a^{-1}.

Although neither utility functions nor their multipliers can be directly compared across agents, relative changes in μ_a can. Suppose that some change in the economy induces a change $\mu_a \to \mu_a + \delta\mu_a$. The corresponding relative changes $\delta\mu_a/\mu_a$ are scale invariant and can be compared across agents. We will later return to this point when we discuss prices.

5. Economic entropy

5.1. The entropy of the whole economy

Substituting the distribution (15) back into (14) yields the maximized value of the information entropy $\mathcal{S}_{\max}$. The result, $\mathcal{S}_{\max} = S(X, Y, U)$ will henceforth be called the economic entropy, or just, the entropy,

$$S(X, Y, U) = \log \zeta(\alpha, \beta, \gamma, \mu) + \sum_g [(\alpha_g - \beta_g)X_g + (\gamma_g - \beta_g)Y_g] + \sum_a \mu_a U_a. \tag{32}$$

The economic entropy can be interpreted as a measure of uncertainty about which agents receive which goods: a higher entropy reflects a more uniform distribution of inputs and outputs over the various agents.

Using Eqs. (19) and (32) the small change in S induced by small changes in the constraints δX, δY, and δU_a is

$$\delta S = \sum_g [(\alpha_g - \beta_g)\delta X_g + (\gamma_g - \beta_g)\delta Y_g] + \sum_a \mu_a \delta U_a, \tag{33}$$

or, in terms of δF_g and δV_g,

$$\delta S = \sum_g (\pi_g \delta F_g + \lambda_g \delta V_g) + \sum_a \mu_a \delta U_a. \tag{34}$$

We can therefore write $S = S(F, V, U)$ where

$$\frac{\partial S}{\partial F_g} = \pi_g, \qquad \frac{\partial S}{\partial V_g} = \lambda_g, \qquad \frac{\partial S}{\partial U_a} = \mu_a. \tag{35}$$

5.2. The signs of π_g and λ_g

Consider a change in the economy in which the production of a particular good g is slightly increased, $\delta F_g > 0$, when all other variables are held constant. In particular, since $V_g = Y_g - X_g$ is held constant the change $\delta F_g > 0$ is achieved by increasing both X_g and Y_g by equal amounts, $\delta Y_g = \delta X_g > 0$. This leads to an increase in entropy, $\delta S > 0$, because of the larger number of ways that the goods $X_g + Y_g = F_g$ can be distributed over the agents. Therefore, $\pi_g = \partial S/\partial F_g > 0$. This is an important ingredient in establishing the interpretation that the π_g's play the role of prices.

Similarly, a change in the economy in which $\delta V_g > 0$ with all else held constant including $F_g = Y_g + X_g$ is achieved by changing X_g and Y_g by equal amounts of opposite sign, $\delta Y_g = -\delta X_g > 0$. This means that the relative mix of input and consumption has changed: a certain amount of good g has been shifted from being used as an input to being used for consumption. One expects that $\delta Y_g > 0$ will increase the entropy while $\delta X_g < 0$ will decrease the entropy. Which one of these contributions dominates depends on the particular economy. Therefore the change in total entropy and the sign of the multiplier $\lambda_g = \partial S/\partial V_g$ can be either positive or negative.

5.3. The entropy of individual agents

Entropy expressions similar to (32) and (33) can be written for the individual agents. Let $X_a = \{X_{ag}\}$, $F_a = \{F_{ag}\}$, and $Y_a = \{Y_{ag}\}$ stand for the sets of variables (28) and (29) associated to the agent *a*. The entropies of the distributions (26) and (27) are

$$\begin{aligned} S_a^x(X_a, F_a) &= -\sum_x P_a(x_a|\alpha, \beta) \log P_a(x_a|\alpha, \beta) \\ &= \log \zeta_a^x(\alpha, \beta) + \sum_g (\alpha_g X_{ag} - \beta_g F_{ag}), \end{aligned} \tag{36}$$

and

$$S_a^y(Y_a, U_a) = -\sum_y P_a(y_a|\gamma, \mu_a) \log P_a(y_a|\gamma, \mu_a) \tag{37}$$

$$= \log \zeta_a^y(\gamma, \mu_a) + \sum_g \gamma_g Y_{ag} + \mu_a U_a. \tag{38}$$

The total entropy for agent a is

$$S_a(X_a, F_a, Y_a, U_a) = S_a^x(X_a, F_a) + S_a^y(Y_a, U_a), \tag{39}$$

and the small entropy change induced by small changes in the constraints δX_{ag}, δY_{ag}, and δU_a is

$$\delta S_a = \delta S_a^x + \delta S_a^y \tag{40}$$

where

$$\delta S_a^x = \sum_g (\alpha_g \delta X_{ag} - \beta_g \delta F_{ag}), \tag{41}$$

$$\delta S_a^y = \sum_g \gamma_g \delta Y_{ag} + \mu_a \delta U_a. \tag{42}$$

Equivalently,

$$\frac{\partial S_a}{\partial X_{ag}} = \alpha_g, \qquad \frac{\partial S_a}{\partial F_{ag}} = -\beta_g, \tag{43}$$

$$\frac{\partial S_a}{\partial Y_{ag}} = \gamma_g, \qquad \frac{\partial S_a}{\partial U_a} = \mu_a. \tag{44}$$

Note that while $F_g = X_g + Y_g$ holds for the economy as a whole, there is no analogous expression for the individual agents: in general $F_{ag} \neq X_{ag} + Y_{ag}$. This is, of course, consistent with economic behavior and with the more traditional way of modeling an economy.

6. Prices are Lagrange multipliers

6.1. The π's are prices

Prices are determined by supply and demand. From the macroscopic point of view the global supply of any good g is represented by the quantity F_g while the demand is ultimately represented through its contribution to the utility functions U_a. Rewrite Eq. (34) as

$$\sum_a \mu_a \delta U_a = -\sum_g (\pi_g \delta F_g + \lambda_g \delta V_g) + \delta S, \tag{45}$$

and consider a small exchange in the supply of two goods: suppose F_1 is increased by a small amount, $F_1 \to F_1 + \delta F_1$, while F_2 is simultaneously decreased, $F_2 \to F_2 - \delta F_2$. Assuming nothing else is changed (*i.e.*, all other δF, δV, and δS are zero) we have

$$\sum_a \mu_a \delta U_a = -\pi_1 \delta F_1 + \pi_2 \delta F_2. \tag{46}$$

Therefore if all $\delta U_a = 0$, then

$$\pi_1 \delta F_1 = \pi_2 \delta F_2. \tag{47}$$

This shows that the economy as a whole is indifferent (as measured by a vanishing change in all U_a) to a trade of an amount δF_2 of good 2 by an amount δF_1 of another good 1, provided the amounts δF_1 and δF_2 are in the inverse ratios of their respective Lagrange multipliers.

On the other hand, if the global prices of these same goods are $\tilde{\pi}_1$ and $\tilde{\pi}_2$ one expects that the condition for indifference be

$$\tilde{\pi}_1 \delta F_1 = \tilde{\pi}_2 \delta F_2. \tag{48}$$

Such relations hold for any pair of goods. This suggests that the equilibrium price $\tilde{\pi}_g$ of good g is proportional to the corresponding multiplier,

$$\tilde{\pi}_g = c\pi_g. \tag{49}$$

Since π_g and $\tilde{\pi}_g$ only differ by an uninteresting scale constant from now on we will also refer to π_g as the prices.

Eq. (47) shows that for this interpretation to make sense it is necessary that all π_gs have the same sign. Our previous discussion in Section 5.2 showed that this is indeed the case, $\pi_g > 0$ and therefore the constant c is positive.

6.2. The budget constraint

All economic modeling must represent the fact that the capacity of individual agents to consume is constrained by their endowment or equivalently, their ability to produce. The connection between production functions f_{ag} and the expected utilities U_a (or equivalently the associated multipliers μ_a) is introduced through the budget constraint. For each agent we require that the expected total cost of consumption matches the expected value of production, Eq. (12),

$$\sum_g \pi_g (F_{ag} - X_{ag}) = \sum_g \pi_g Y_{ag}. \tag{50}$$

The expectations on the left hand side, Eq. (28), depend on production functions; those on the right hand side, Eq. (29), depend on the unknown expected utility U_a through the multiplier μ_a. Imposing this budget constraint allows us to calculate μ_a in terms of the agent's production functions and the other parameters (the Lagrange multipliers α_g, β_g, and γ_g, or equivalently, π_g and λ_g) that describe the economy as a whole. Except in very simple cases the calculation of μ_a will have to be carried out numerically.

The relationship between prices π_g and multipliers μ_a (or, equivalently, the expected utilities U_a) is one of self-consistency: Given the multipliers μ_a (which regulate the expected utilities U_a) and α_g, β_g, and γ_g (which regulate supply F_g and demand Y_g) the prices π_g are determined by Eq. (23). But the U_as themselves are unknown; to determine them the additional information provided by the budget constraints is needed: Given the prices π_g, the multipliers μ_a are determined by Eq. (50). Thus, in the entropic framework, prices and expected utilities are determined endogenously—they are endogenous to the economy.

To summarize, in this entropic framework prices are essentially the Lagrange multipliers that are endogenously determined. They convey all of the available information about all of the economic activities and agents' behavior. This is consistent with the traditional economic modeling, but here we can explicitly see how each Lagrange multiplier represents the informational impact of the corresponding constraint on the maximal level of entropy—also a natural informational quantity.

7. Two economies in equilibrium

Consider two economies A and B that are maintained separated by some trade barrier. The trade barrier is lifted and the combined economy C reaches a new equilibrium. The initial equilibrium situation is described by variables $F_{Cg}^{(\mathrm{in})} = F_{Ag}^{(\mathrm{in})} + F_{Bg}^{(\mathrm{in})}$, $V_{Cg}^{(\mathrm{in})} = V_{Ag}^{(\mathrm{in})} + V_{Bg}^{(\mathrm{in})}$, $U_a^{(\mathrm{in})}$, and $U_b^{(\mathrm{in})}$ ($a \in A$ and $b \in B$ refer to agents in either sub-economy A or B and "in" refers to "initial"). Similarly the new equilibrium situation for the joint economy is described by variables F_{Cg}, V_{Cg}, and U_c (c ranges over all agents $a \in A$ and $b \in B$). The values of these variables need not coincide with the initial values—in general there is no reason why any of these quantities should be conserved. What can we tell about the new equilibrium state?

It seems reasonable to conjecture that the removal of a barrier will probably lead to an increase in entropy, $S_{\mathrm{final}} \geq S_{\mathrm{initial}}$, which would provide an economic analogue to the second law of thermodynamics. However, as shown in Ref. [15] the derivation of the 2nd law of thermodynamics relies on explicit knowledge that the dynamics describing the approach to equilibrium is a Hamiltonian. In the absence of further knowledge of the dynamics that governs the approach to equilibrium the universal increase of *economic* entropy must remain a plausible conjecture.

7.1. Prices and equilibrium

In the absence of some additional information about the dynamical process leading to the new equilibrium the quantities F_{Cg}, V_{Cg} and U_c will remain unknown. However, we do know that the new equilibrium is a maximum entropy state and according to the completeness assumption (Section 4.1) we also know that F_{Cg}, V_{Cg} and U_c capture the relevant information that ought to be known to make a reliable inference. We therefore proceed as if these values were known which allows us to make some statements about the redistribution (reallocation) of resources and goods among the two economies.

Our goal is to find the new values F_{Ag}, V_{Ag}, F_{Bg}, and V_{Bg} that maximize the joint final entropy for some fixed (but perhaps unknown) totals $F_{Cg} = F_{Ag} + F_{Bg}$ and $V_{Cg} = V_{Ag} + V_{Bg}$. Since in equilibrium the agents $a \in A$ and $b \in B$ are statistically independent the corresponding entropies are additive

$$S_C(F_T, V_T, U_c) = S_A(F_A, V_A, U_a) + S_B(F_T - F_A, V_T - V_A, U_b). \tag{51}$$

For any given good g the maximum with respect to variations of F_{Ag} is attained when

$$0 = \frac{\partial S_C}{\partial F_{Ag}} = \frac{\partial S_A}{\partial F_{Ag}} + \frac{\partial S_B}{\partial F_{Ag}}, \tag{52}$$

which, can be written as

$$\frac{\partial S_A}{\partial F_{Ag}} = \frac{\partial S_B}{\partial F_{Bg}}. \tag{53}$$

Similarly the maximum with respect to variations of V_A is attained when

$$0 = \frac{\partial S_T}{\partial V_{Ag}} = \frac{\partial S_A}{\partial V_{Ag}} + \frac{\partial S_B}{\partial V_{Ag}}, \tag{54}$$

or

$$\frac{\partial S_A}{\partial V_{Ag}} = \frac{\partial S_B}{\partial V_{Bg}}. \tag{55}$$

Using Eq. (35), $\partial S/\partial F = \pi$ and $\partial S/\partial V = \lambda$, we see that the maximum entropy condition for the new equilibrium macrostate is that the Lagrange multipliers be equal,

$$\pi_{Ag} = \pi_{Bg} \quad \text{and} \quad \lambda_{Ag} = \lambda_{Bg}. \tag{56}$$

Thus, the principle of maximum entropy leads to the prediction that the new state of statistical equilibrium is achieved when every π_g and λ_g attain equal values across the two economies. This, together with the earlier discussion in Section 5 (that $\pi_g > 0$) and Section 6, provides further evidence in favor of our earlier interpretation of the multipliers π_g as prices. Concerning the λ's, just as the prices π_g regulate the flow of goods, the λ_g regulate the input/consumption mix. It is interesting to note that these two sets of conditions play the same role for statistical equilibrium as the vanishing of the aggregate excess demand functions play in the traditional Walrasian equilibrium.

7.2. An open economy: a small economy interacting with a larger one

This is a special case of the two-economy case. The two economies A and B interact, goods flow in the appropriate directions until a new equilibrium is reached. The interesting element here is that if economy B is much smaller than economy A then the latter will hardly be affected by the interaction. The actual process is one in which as the interaction proceeds the multipliers π_{Ag} and λ_{Ag} remain fixed at their initial values while π_{Bg} and λ_{Bg} will evolve until they match the A values. The larger economy acts as a reservoir capable of exchanging essentially arbitrary amounts of goods and resources with the smaller economy without itself being affected, that is, the values of its own Lagrange multipliers remain fixed.

7.3. Stability of the economic macrostate

Consider an economy C composed of two sub-economies A and B. In order for the macrostate to represent a situation of stable equilibrium the total entropy must be a maximum and not just stationary. Therefore (omitting all superfluous variables) at the maximum the first derivative of

$$S_C(F_C) = S_A(F_A) + S_B(F_T - F_A) \tag{57}$$

must vanish, Eq. (52), and its second derivative must be negative,

$$\frac{\partial S_C}{\partial F_{Ag}} = 0 \quad \text{and} \quad \frac{\partial^2 S_C}{\partial F_{Ag}^2} < 0. \tag{58}$$

Therefore,

$$\frac{\partial^2 S_C}{\partial F_{Ag}^2} = \frac{\partial^2 S_A}{\partial F_{Ag}^2} + \frac{\partial^2 S_B}{\partial F_{Bg}^2} < 0, \tag{59}$$

and using $\partial S/\partial F_g = \pi_g$ we get

$$\frac{\partial \pi_{Ag}}{\partial F_{Ag}} + \frac{\partial \pi_{Bg}}{\partial F_{Bg}} < 0. \tag{60}$$

In the particular case that the two sub-economies happen to be identical to each other then these two terms are the same. Therefore the condition for stability is that

$$\frac{\partial \pi_g}{\partial F_g} < 0. \tag{61}$$

Thus, the condition for stability is that an increase in the total availability of any good g must lead to a decrease in its price. The stability condition is crucial for the interpretation of π_g as the price of g: since goods tend to flow from low to high values of π_g, a stable economy can only be achieved if the result of any spontaneous fluctuation in π_g is to induce a flow that returns the system to uniform prices. This is the economic analogue of Le Chatelier's principle.

A similar argument can be carried out for the multiplier λ and the resulting condition for stability is

$$\frac{\partial \lambda_g}{\partial V_g} < 0. \tag{62}$$

A stable equilibrium requires that the economy be stable with respect to the flow of goods and also with respect to the input/consumption mix. Therefore both conditions (61) and (62) must be satisfied for every good g.

8. An example

To illustrate the ideas above it is useful to study in some detail a simple idealized model where calculations can be carried out explicitly. The model economy is such that production and utility functions are linear,

$$f_{ag'} = f_{ag'}(x_{a1g'}, x_{a2g'} \ldots x_{aGg'}) = \sum_g b_{agg'} x_{agg'}, \tag{63}$$

and

$$u_a = u_a(y_{a1}, y_{a2} \ldots y_{aG}) = \sum_g c_{ag} y_{ag}. \tag{64}$$

Recall that $x_{agg'}$ is the amount of g used by a to produce g' and $x_{ag} = \sum_{g'} x_{agg'}$ is the total amount of g used by a. Here $b_{agg'}$ and c_{ag} are the appropriate constant coefficients for each agent and good.

In the standard microeconomic theory [16] linear production functions and linear utility functions are not useful—they may lead to trivial equilibria. In the entropic framework, however, linear functions are useful because neither is meant to be maximized. They perform the desired role of differentiating different agents according to productivity and different goods according to preference. Most importantly, while admittedly extremely idealized this simple model already shows highly non-trivial and non-linear behavior.[8]

The macrostate $P(x, y)$ is given by Eqs. (25), (26), and (27). For the linear model the distributions of $x_a = \{x_{agg'}\}_a$ and $y_a = \{y_{ag}\}_a$ for a given agent a are

$$P_a(x_a|\alpha, \beta) = \frac{1}{\zeta_a^x} \exp\left[-\sum_{gg'} (\alpha_g - \beta_{g'} b_{agg'}) x_{agg'}\right] \tag{65}$$

and

$$P_a(y_a|\gamma, \mu_a) = \frac{1}{\zeta_a^y} \exp\left[-\sum_g (\gamma_g + \mu_a c_{ag}) y_{ag}\right]. \tag{66}$$

Thus, conditional on the Lagrange multipliers, the inputs $x_{agg'}$ are statistically independent,

$$P_a(x_a|\alpha, \beta) = \prod_{gg'} P_{agg'}(x_{agg'}|\alpha_g, \beta_{g'}) \tag{67}$$

where

$$P_{agg'}(x_{agg'}|\alpha_g, \beta_{g'}) = \frac{1}{\zeta_{agg'}^x} \exp[-\xi_{agg'} x_{agg'}] \quad \text{where } \xi_{agg'} = \alpha_g - \beta_{g'} b_{agg'}. \tag{68}$$

The same is true for the consumption goods: conditional on the Lagrange multipliers, the variables y_{ag} are statistically independent,

$$P_a(y_a|\gamma, \mu_a) = \prod_g P_{ag}(y_{ag}|\gamma_g, \mu_a) \tag{69}$$

where

$$P_{ag}(y_{ag}|\gamma_g, \mu_a) = \frac{1}{\zeta_{ag}^y} \exp[-\chi_{ag} y_{ag}] \quad \text{where } \chi_{ag} = \gamma_g + \mu_a c_{ag}. \tag{70}$$

8.1. *Expected values: the Bose–Einstein distribution*

We now analyze the distributions above. Consider inputs first. The partition function in Eq. (68) is a geometric series, therefore

$$\zeta_{agg'}^x = \sum_{z=0}^{\infty} e^{-\xi_{agg'} z} = \frac{1}{1 - e^{-\xi_{agg'}}}, \tag{71}$$

provided the series converges. The condition for convergence is

$$e^{-\xi_{agg'}} < 1 \quad \text{or} \quad \xi_{agg'} = \alpha_g - \beta_{g'} b_{agg'} > 0. \tag{72}$$

This result turns out to be significant and will be discussed further below. The main point is that values of α_g and $\beta_{g'}$ such that $\xi_{agg'}$ is negative are ruled out as unattainable because the three constraints on X_g, $F_{g'}$, and normalization are incompatible.

8 To our knowledge the only model that makes a similar use of linear functions is [8] which deals with the special case of a production economy. The present framework is more general.

The expected amount of input g accessible to agent a for production of good g' is

$$X_{agg'} = \langle x_{agg'} \rangle = -\frac{\partial \log \zeta^x_{agg'}}{\partial \xi_{agg'}} = \frac{1}{e^{\xi_{agg'}} - 1}, \tag{73}$$

which is characteristic of a Bose–Einstein distribution.[9] The total amount of resource g used by agent a is

$$X_{ag} = \sum_{g'} X_{agg'} = \sum_{g'} \frac{1}{e^{\xi_{agg'}} - 1}. \tag{74}$$

The amount of good g' produced by agent a is

$$F_{ag'} = \sum_{g} b_{agg'} X_{agg'} = \sum_{g} \frac{b_{agg'}}{e^{\xi_{agg'}} - 1}. \tag{75}$$

A similar analysis applies to consumption goods. The partition function in Eq. (70) is

$$\zeta^y_{ag} = \sum_{z=0}^{\infty} e^{-\chi_{ag} z} = \frac{1}{1 - e^{-\chi_{ag}}}, \tag{76}$$

provided

$$\chi_{ag} = \gamma_g + \mu_a c_{ag} > 0. \tag{77}$$

The expected amount of consumer good g accessible to agent a is

$$Y_{ag} = \langle y_{ag} \rangle = -\frac{\partial \log \zeta^y_{ag}}{\partial \chi_{ag}} = \frac{1}{e^{\chi_{ag}} - 1}, \tag{78}$$

which is also a Bose–Einstein distribution. The expected utility of agent a is

$$U_a = \sum_{g} c_{ag} Y_{ag} = \sum_{g} \frac{c_{ag}}{e^{\chi_{ag}} - 1}. \tag{79}$$

8.2. Budget constraints

Each agent a is subject to a budget constraint, Eq. (50). Given the state of the economy as specified by the values of α_g, β_g, and γ_g for each good g we can calculate the expected budget B_a available to an agent a,

$$\begin{aligned} B_a &= \sum_{g} \pi_g F_{ag} - \sum_{g} \pi_g X_{ag} \\ &= \sum_{gg'} \frac{\pi_{g'} b_{agg'} - \pi_g}{\exp(\alpha_g - \beta_{g'} b_{agg'}) - 1} \end{aligned} \tag{80}$$

where $\pi_g = \frac{1}{2}(\alpha_g + \gamma_g) - \beta_g$. The corresponding expected expenditure is

$$B_a = \sum_{g} \pi_g Y_{ag} = \sum_{g} \frac{\pi_g}{\exp(\gamma_g + \mu_a c_{ag}) - 1} \tag{81}$$

which is an implicit equation for the multiplier μ_a of each agent.

Eq. (80) shows that an agent that manages to increase its productivity (say, by increasing one of the productivity constants $b_{agg'}$) will have access to a bigger budget B_a. Eq. (81) then shows that, other things being equal, the bigger budget B_a leads to a lower value of μ_a. The agent is effectively wealthier because, as we saw in Section 4.3, the lower value of μ_a implies a higher probability of attaining microstates of higher utility.

8.3. Discussion: monopoly vs. competitive market

Eqs. (72) and (73) show that as $\xi_{agg'} \to 0$ the expected amount of g used by agent a diverges. A more careful calculation shows that a macroscopic fraction of the total amount of good g available to the whole economy is concentrated into a single agent a (or a small group of agents) characterized by the particular constant $b_{agg'}$ that leads to $\xi_{agg'} = 0^+$. Indeed,

$$X_{agg'} = \frac{1}{e^{\xi_{agg'}} - 1} \approx \frac{1}{\xi_{agg'}} = \frac{1}{\alpha_g - \beta_{g'} b_{agg'}}. \tag{82}$$

[9] The Bose–Einstein distribution for linear technologies was also derived in Ref. [8].

The fact that the value $X_{agg'}$ cannot be larger than the total amount X_g available to the economy as a whole places a lower bound on how small the value of $\xi_{agg'}$ can be.

The production of good g' is correspondingly increased,

$$F_{ag'} = \sum_g \frac{b_{agg'}}{e^{\xi_{agg'}} - 1} \approx \frac{b_{agg'}}{\xi_{agg'}}. \tag{83}$$

The agent a has an effective market power over both the input/resource g and the produced good g'. (Agent a is both a monopsony and a monopoly.)

We can also see from Eq. (80) that the same condition, $\xi_{agg'} \to 0^+$, leads to

$$B_a \approx \frac{\pi_{g'} b_{agg'} - \pi_g}{\xi_{agg'}} \tag{84}$$

which means that agent a achieves an immense amount of wealth.

It could be that a group of several agents is characterized by the particular productivity constant $b_{agg'}$. Then the situation is not a monopoly, but the market will still be very concentrated. At the other extreme where all agents are roughly equally productive one expects a more uniform distribution and a competitive market.

A similar effect can occur with consumption. As $\chi_{ag} \to 0$ Eq. (78) shows that a macroscopic fraction of a particular good g may end up being consumed by just one or a few agents. Whether this happens, see Eq. (77), depends on the value of μ_a, which is determined by the budget constraint, and on the relative signs of γ_g and μ_a.

This phenomenon where a macroscopic fraction of resources or wealth accumulates into a single or a few agents has a physics analogue when a macroscopic fraction of all molecules in a system condense into the ground state. This is known as the Bose–Einstein condensation. The study of the detailed behavior of this special regime can be carried out within our entropic framework and is interesting for its own sake. However, it lies outside the main objective of this paper which is to propose a general framework for modeling.

9. Summary and discussion

We have presented an entropic, information-theoretic framework for modeling economies. The approach is of fairly broad applicability and can be easily adapted to a wide variety of economic systems. Our framework does not compete with the more traditional approaches; instead, it complements them.

A main difference between our approach and the more traditional ones is that we model the economy from the point of view of an external observer who has limited macroscopic information about the system. This means that we deal with an inference problem that is inherently under-determined. The solution, provided by the method of maximum entropy, takes account of relevant macroscopic information we have while being maximally non-committal about all those microscopic details that we do not know or do not care about. The result is the probability distribution of resources (endowments) and goods across all agents in the economy. Such a macrostate describes a situation of statistical economic equilibrium. It allows us to investigate the conditions under which the distribution of wealth is relatively uniform as well as all the different types of markets (from competitive to monopoly and/or monopsony). The concept of economic entropy is instrumental in providing us with the required measure of uniformity of the distribution (or the marginal distributions of interest).

In the entropic framework prices arise naturally as statistical constructs—they are the Lagrange multipliers. Just as in traditional economic models they convey the relevant information about equilibrium and they regulate the flow of inputs and outputs as the economy opens for trade. Furthermore, they also capture the informational content of each one of the model's constraints.

A basic feature of this approach is that it is easy to generalize our model by incorporating additional information through additional constraints. The analysis of any such generalized models, which will necessarily have to be carried out numerically, will enable us to investigate the meaning and impact of each additional piece of information and eventually provide us with the necessary tools for investigating different policy and behavioral scenarios.

Another interesting direction on which our framework could be extended is to account for varying numbers and types of agents—how agents are created and destroyed and how they might move from one economy to another. In this respect, we might remark on the approach by Yakovenko et al. [12,13] which is clearly complementary to ours. In physics one can study the problem of how a collection of atoms are distributed over given energy levels. One can also study the related but different problem of how a certain amount of energy is shared among systems in thermal contact. The analysis by Yakovenko et al. is analogous to the former problem, ours to the latter: Yakovenko et al. study how agents are distributed over different money or income levels. In contrast, we study how goods and resources are shared over agents in economic equilibrium. One advantage of Yakovenko's approach is the ease with which varying numbers of essentially identical agents are included in the analysis. One advantage of our approach is the ease with which we account for the multitude of differences among agents—different production functions, different preferences, etc.

Acknowledgment

A.C. gratefully acknowledges the generous hospitality of the Info-Metrics Institute.

References

[1] R.D. Rosenkrantz (Ed.), E.T. Jaynes: Papers on Probability, Statistics and Statistical Physics, Reidel, Dordrecht, 1983. Papers online at http://bayes.wustl.edu.
[2] E.T. Jaynes, in: G.L. Bretthorst (Ed.), Probability Theory: The Logic of Science, Cambridge UP, 2003.
[3] A. Golan, Information and entropy econometrics: a review and synthesis, Found. Trends Econom. 2 (1–2) (2008) 1–145.
[4] A. Caticha, Entropic inference and the foundations of physics, in: Monograph Commissioned by the 11th Brazilian Meeting on Bayesian Statistics—EBEB-2012, USP Press, São Paulo, 2012. http://www.albany.edu/physics/ACaticha-EIFP-book.pdf.
[5] H. Theil, Economics and Information Theory, North Holland, Amsterdam, 1967.
[6] N. Georgesçu-Roegen, The Entropy Law and the Economic Process, Harvard U. P., Cambridge, 1971.
[7] E. Smith, D.K. Foley, Classical thermodynamics and economic general equilibrium, J. Econom. Dynam. Control 32 (2008) 7–65.
[8] A. Golan, A multivariable stochastic theory of size distribution of firms with empirical evidence, Adv. Econom. 10 (1994) 1–46.
[9] D.K. Foley, A statistical equilibrium theory of markets, J. Econom. Theory 62 (1994) 321.
[10] E. Smith, D.K. Foley, B.H. Good, Unhedgeable shocks and statistical economic equilibrium, Econom. Theory 52 (2013) 187.
[11] A.A. Toda, Existence of a statistical equilibrium for an economy with endogenous offer sets, Econom. Theory 45 (2010) 379.
[12] V.M. Yakovenko, J.B. Rosser Jr., Statistical mechanics of money, wealth, income, Rev. Modern Phys. 81 (2009) 1703.
[13] A. Banerjee, V.M. Yakovenko, Universal patterns of inequality, New J. Phys. 12 (2010) 075032.
[14] A. Golan, J.W. Bono, Identifying strategies and beliefs without rationality assumptions, http://dx.doi.org/10.2139/ssrn.1611216.
[15] E.T. Jaynes, Gibbs vs. Boltzmann entropies, Amer. J. Phys. 33 (1965) 391.
[16] H.R. Varian, Microeconomic Analysis, Norton, New York, 1992.