



The Google thought experiment: rationality, information and equilibrium in an exchange economy

Sylvain Barde

► To cite this version:

Sylvain Barde. The Google thought experiment: rationality, information and equilibrium in an exchange economy. 2009. hal-01069373

HAL Id: hal-01069373

<https://sciencespo.hal.science/hal-01069373>

Preprint submitted on 29 Sep 2014

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

**THE GOOGLE THOUGHT EXPERIMENT: RATIONALITY,
INFORMATION AND EQUILIBRIUM IN AN EXCHANGE ECONOMY**

OFCE

N° 2009-34

DECEMBER 2009

Sylvain Barde
OFCE

The Google thought experiment: rationality, information and equilibrium in an exchange economy*

Sylvain Barde^a

^a*OFCE - Department for Research on Innovation and Competition, 250 rue Albert Einstein, Valbonne - Sophia Antipolis, 06560 France. tel : +33 (0)4 93 95 42 38, email: sylvain.barde@ofce.sciences-po.fr*

November 2009

Abstract

Following Becker (1962), an information-theoretical thought experiment is developed to investigate whether the equilibrium properties of an exchange economy depend on the rational behaviour of agents. Transactions are logged through a communication channel into an external observer's dataset, represented by Google. At some point this data connection fails and Google no longer receives the updates encoding the transactions.

It is shown that Google can nevertheless make sharp predictions concerning the state of the economy. In particular, a stable long run distribution of endowments is expected, as well as a set of price-like variables. By construction this prediction does not rest on the rationality of agents, because the information-theoretical setting forces Google to treat the missed updates as random variables.

JEL classification: C02, C62, D51, D80.

Keywords: Information theoretical measure, maximum entropy principle, entropy concentration theorem.

1 Introduction

In a recent contribution, Smith and Foley (2008) review the long history of analogies between economics and classical thermodynamics. In particular, they point out the historical link between the original cardinal approach to utility and the use of gradients on the one hand and the physical concepts of energy and forces on the other. They further

*The usual disclaimers apply.

explain that as economic theory has matured, it has distanced itself from such comparisons or analogies. Nevertheless it will be shown, in agreement with Smith and Foley, that “there are important insights still to be gleaned from considering the relation of neoclassical economics to classical thermodynamics” (Smith and Foley, 2008, p. 9), in particular with respect to the problem raised by Becker (1962) regarding the importance of the rationality of agents in the market allocation process. In order to do so, however, classical thermodynamic analogies will be replaced by the information-theoretical definition of entropy introduced by Shannon (1948), and will be investigated in a thought experiment by using the maximum entropy principle of Jaynes (1957a,b).

The setup of the thought experiment is simple, and puts Google in the position of the subjective Hayekian observer of the economy: Google is assumed to initially possess a dataset that describes the detailed endowments of all the agents in an exchange economy.¹ A data connection allows this dataset to be updated every time a transaction occurs i.e. the incoming updates change the endowment levels in the relevant entries. A real-life example of such a setup would be a real-time back-up of a stock exchange book, where bids or transactions between agents are logged as part of the market-making process. From the point of view of Google, the updating process can be described with a Markov process, where each update reveals the detailed endowment state that the economy has just transitioned to. At a given point in time the data connection fails: although normal market activity continues, Google’s dataset is no longer updated. At the time of the failure, the dataset fully describes the state of the economy. However, one intuitively understands that as time passes by the usefulness of this data in describing the state of the economy will decline, leaving Google in a position of subjective ignorance described in Hayek (1945).

This thought experiment raises several questions. Analysing the allocation problem using the maximum entropy method suggested by Jaynes (1957a,b), one can show that although Google becomes completely ignorant of the detailed state of the economy, is it still possible for it to make precise predictions about its overall state. Furthermore, these

¹The motivation behind this choice is that Google is indeed an observer of agent transactions, and as such provides a good illustration of the information-theoretical measure used here.

predictions are consistent with the general properties of economic equilibrium, namely the existence of stable endowment distributions and market prices that equalise in the presence of trade. However, because the setup of the thought experiment forces Google to treat the updates as random variables in a Markov process, these predicted properties do not rely on any assumptions regarding underlying the rationality or behaviour of agents, which tends to confirm the suggestion made in Becker (1962) and Gode and Sunder (1945) that rationality is not the only source of market regularities.

The rest of the paper is structured as follows. Section 2 describes the updating process and the information theoretical methodology the observer can use to predict the overall state of the economy. Section 3 then describes the properties of the prediction, in particular the existence of a stable distribution and the existence of price-like variables. Section 4 discusses the implications of this results, particularly with respect to economic ontologies and finally section 5 concludes.

2 The Google thought experiment

2.1 Notation and definitions

There are N agents and K markets in the economy. For all $k \in K$, the economy-wide endowment is $\omega^k \in \mathbb{N}$. This allows us to define on each market k a set of endowment levels $\{\varepsilon^k\}$ that agents can possess, with ε_i^k representing the i^{th} endowment level in the set.

$$\varepsilon_i^k \in \{0, 1, 2, \dots, \omega^k\} \tag{1}$$

A detailed description of the state of the economy consists of a dataset with $N \times K$ entries, each of them containing the n^{th} agent's endowment level of the k^{th} good. Within this setup a *record* is a row vector that consists of the K entries that fully describe n^{th} agent's endowment. It is shown below that Google's uncertainty as to the state of the economy at time t is related to his uncertainty as to the average information content of a

record.

The frequency of an endowment level p_i^k , which is also the probability that a randomly selected agent has that level of endowment, is simply defined as n_i^k , the number of agents with endowment ε_i^k , over the total population N .

$$p_i^k = \frac{n_i^k}{N}$$

The number of ways the N records of the k^{th} market can result in the probability distribution p_i^k over the endowment level is given by the following multiplicity:

$$W(p_i^k) = \frac{N!}{\prod_{i=1}^{\omega^k} N p_i^k!} \quad (2)$$

Furthermore, given the K markets, we will assume that the overall multiplicity of the states is the product of the multiplicities on each market. This implicitly means that there is no interaction between the markets, i.e. that the endowment level of a given good does not condition the endowment level of other goods. While this might seem like a strong assumption, this will be shown to be consistent with the thought experiment set up in the next section.

$$W = \prod_{k=1}^K W(p_i^k)$$

As shown by Foley (1994), the standard definition for entropy is the log of this multiplicity.

$$S = \sum_{k=1}^K \ln W(p_i^k) \quad (3)$$

Appendix A.1 shows that the link with the information entropy of the probability distribution is given by:

$$S = N \sum_{k=1}^K H(p_i^k) \quad (4)$$

Where $H(p_i^k)$ is the Shannon (1948) entropy of the distribution, i.e. the expected information content of a message informing us of the outcome of a draw on this distribution.

$$H(p_i^k) = - \sum_{i=1}^{\omega^k} p_i^k \ln p_i^k \quad (5)$$

Equation (4) shows the relation between the two interpretations of entropy, the objective S which represents the log of the multiplicity of outcomes, and the subjective H which represents the uncertainty or average information content of a record that the observer is ignorant about.²

The total entropy of the system is therefore the information entropy per record, multiplied by the number of records in the dataset. This implicitly means that the expected information content is the same for all unknown records, which is consistent with the setup of thought experiment, as a different result would imply that Google still possesses information about a specific agent.

2.2 Uncertainty in the updating process

In order to describe the increase in uncertainty, i.e. establish the expected rate at which the data contained in Google’s $N \times K$ dataset decays, two elements need to be distinguished: the first is the information content of the average entry, which will quantify the information loss that occurs when an entry is overwritten by an update and the second is the rate at which distinct dataset entries are updated.

Regarding the first point, dataset entries have different information content depending on which of the K markets they describe. Indeed, if the k^{th} good is rare, the set of endowment levels (1) is limited, and the associated probabilities are relatively high. Therefore, the information loss (5) associated with a missed update on that market is low: learning that an agent has changed his endowment level does not create much uncertainty, as there are not many alternative possibilities to guess from. On the other hand, if the endowment is large there are many endowment levels available for each agent, thus

²Because the natural logarithm is used here, the information entropy of a record is measured in “nats”, where a nat is equal to $1/\ln 2 \approx 1.44$ bits of information.

loosing information about an entry creates much more uncertainty.

As a result Google's uncertainty about the content of a missing update has to integrate not only the uncertainty about the possible endowment levels on each market, but also the uncertainty as to which of the K markets the update modifies. In the absence of any knowledge about the updating process, Google has to assume that records concerning the different markets are updated with equal probability $p^K = 1/K$.³ Therefore, the joint probability that a missed update shifts a record's k^{th} endowment to ε_i^k is given by $q_i^k = p^K \times p_i^k$.

Based on the standard definition of conditional entropy given in Theil (1967), the expected information content of an update U conditional on the selection of a market K is:

$$H(U|K) = - \sum_{k=1}^K \sum_{i=1}^{\omega^k} q_i^k \ln \frac{q_i^k}{p^K}$$

Replacing the probabilities q_i^k and rearranging gives the expected information content per updated entry, which is simply the average of the expected information for each of the K entries.

$$H(U|K) = -\frac{1}{K} \sum_{k=1}^K \sum_{i=1}^{\omega^k} p_i^k \ln p_i^k \quad (6)$$

The second element to consider is the rate at which the information existing dataset decays. Although each update carries the same expected information content (6), not all contribute to increasing uncertainty as to the state of the economy: there is an increase in uncertainty, in other words a net information loss, only in the case when a missed update with positive entropy overwrites a known zero-entropy entry. In the thought experiment, this only happens at the first update of an entry. Subsequent updates to the same entry may replace the objective information content, but in terms of uncertainty, Google is as unsure of the information content prior to the following updates as he is

³This is equivalent to the assumption made above regarding the independence of the endowments levels on each market.

after.⁴ Therefore, from Google’s point of view, information is only lost the first time an entry is updated. Given the entropy of an updated entry (6) and the fact that the dataset contains $N \times K$ entries, the total information loss incurred once all entries have been overwritten is indeed equal to the overall amount of information required to describe the dataset given by equation (4).

The expected number of missed updates required for all $N \times K$ entries to have been updated at least once is analogous to the standard coupon collector problem, described for example in Ross (2002). The probability of overwriting the j^{th} entry given that $j - 1$ distinct entries have already been over-written is:

$$p_j = \frac{NK - j + 1}{NK}$$

Hence, the a new entry is overwritten by the first update with probability one, but once $NK - 1$ entries have been overwritten, the last entry is overwritten by an update with probability $1/NK$, which is intuitive: the first few updates that the observer misses will practically all concern distinct entries in the dataset and result in loss of information. The expected amount of updates T required to loose all the initial NK entries is therefore:⁵

$$E(T) = NK \ln NK + \gamma NK \tag{7}$$

If we assume that Google knows the expected rate of arrival of the updates, then it can work out the expected amount of time required for the dataset to decay completely.

2.3 Google’s best prediction under uncertainty

After the T updates have occurred, Google’s best guess of the distribution of endowments over the N agents is the one that maximises the dataset entropy. Indeed, any prediction made about the overall state of the economy at that point must be consistent with the

⁴This possible because the economy-wide endowments ω^k are fixed. This would change in the context of a production economy, as the ω^k endowments will change over time, changing the probability distributions over the endowment sets (1)

⁵ γ is the Euler–Mascheroni constant, which is approximately 0.577.

loss of information measured by the entropy function: Were Google able to predict a distribution with an entropy below the maximum information entropy, this would imply that it still posses some information about the system.

However, as explained in Jaynes (1957a,b) this uncertainty about the state of the system can be bounded by supplying information on known constraints. A first example, which exists by construction, is that the endowment probabilities must sum to one:

$$\forall k, \sum_{i=1}^{\omega^k} p_i^k = 1 \quad (8)$$

In the case of an exchange economy, an additional constraint is that the economy-wide endowments, the ω^k 's, are conserved in the exchanges, with no creation *ex nihilo* or destruction of goods.⁶ In terms of subjective probabilities on the state of a given agent's record, the constraint is that the expected endowment level is equal to the per-capita amount of the good for each market.

$$\forall k, \sum_{i=1}^{\omega^k} p_i^k \varepsilon_i^k = \frac{\omega^k}{N} \quad (9)$$

As outlined in Jaynes (1957a,b), further constraints can be integrated, based on information that Google might have of the mechanisms of the economy or on the availability of partial data. Although (8) and (9) are the minimum number of constraints required to describe an exchange economy, this possibility makes the methodology very general, as an arbitrary number of constraints can be integrated.

Google's best prediction therefore maximises the information entropy of the dataset (5), subject to the known constraints (8) and (9).

$$L = - \sum_{k=1}^K \sum_{i=1}^{\omega^k} p_i^k \ln p_i^k - \sum_{k=1}^K \left(\alpha^k \left(\sum_{i=1}^{\omega^k} p_i^k - 1 \right) + \beta^k \left(\sum_{i=1}^{\omega^k} p_i^k \varepsilon_i^k - \frac{\omega^k}{N} \right) \right) \quad (10)$$

Differentiating with respect to all k probability distributions, one obtains the following

⁶This condition must be modified in the case of a production economy, as goods will be both produced and consumed.

first order conditions:

$$p_i^k = e^{-(\alpha^k+1)} e^{-\beta^k \varepsilon_i^k} \quad (11)$$

Replacing this into (8) allows us to determine the values of the first set of lagrangian multipliers α^k :

$$e^{-(\alpha^k+1)} = \frac{1}{\sum_{i=1}^{\omega^k} e^{-\beta^k \varepsilon_i^k}}$$

In line with the typical notation, we define the following partition function Z^k , which normalises the probability distribution and guarantees that (8) holds.

$$Z^k = \sum_{i=1}^{\omega^k} e^{-\beta^k \varepsilon_i^k}$$

In the typical record, the probability that the k^{th} entry holds the i^{th} endowment is therefore:

$$p_i^k = \frac{e^{-\beta^k \varepsilon_i^k}}{Z^k} \quad (12)$$

By replacing this into (9) one can show that the β^k terms are functions of their respective ω^k only. The value of each market endowment ω^k therefore uniquely defines each β^k parameter.

$$\frac{\sum_{i=1}^{\omega^k} \varepsilon_i^k e^{-\beta^k \varepsilon_i^k}}{Z^k} = \frac{\omega^k}{N}$$

Given the functional form, it is difficult to invert this expression to express β^k as a function of ω^k . However, one can nevertheless obtain from this equation the limit values of β^k given the endowment level:

$$\lim_{\omega^k \rightarrow 0} \beta^k = \infty \quad \lim_{\omega^k \rightarrow \infty} \beta^k = 0$$

If the value of the endowment is zero, then all the probabilities associated with $\varepsilon_i^k > 0$ must be zero, and the probability associated with $\varepsilon_i^k = 0$ must be equal to one. For the second limit, if the good exists in limitless quantity then there is in fact no endowment constraint on the probability distribution, and the entropy maximising probability distribution over the endowment set is uniform. This means that equation (12) must be independent of ε_i^k , which requires β^k to be zero.

3 Properties of Google's best prediction

3.1 The entropy concentration theorem

Google might wonder at this point how likely it is that the predicted distribution (12) is realised relative to alternative distributions. While the distribution that maximises entropy is the most probable one, this does not tell us how probable it is *relative to others*. This is clarified by Jaynes (1983) using the entropy concentration theorem, which not only provides the core justification for the use of this methodology, but also provides a statistical test that can be used for inference: the maximum entropy prediction completely dominates the state space, and is therefore the only rigorous prediction in the absence of detailed information.

The starting point of this analysis is that given two distributions p_i^k and q_i^k , the number of ways p_i^k can be realised over N records relative to the way q_i^k is realised is simply the ratio of their multiplicities (2):

$$\frac{W(p_i^k)}{W(q_i^k)} = \prod_{i=1}^{\omega^k} \frac{N q_i^k!}{N p_i^k!} \quad (13)$$

Using Stirling's approximation and the definition of the information entropy of a record (5), one can express this multiplicity ratio as:⁷

$$\frac{W(p_i^k)}{W(q_i^k)} = \sqrt{\prod_{i=1}^{\omega^k} \frac{q_i^k}{p_i^k}} e^{N(H(p_i^k) - H(q_i^k))} \quad (14)$$

⁷Details on this are given in appendix A.3

Let p_i^k be the entropy-maximising distribution, and let the Google choose a fixed multiplicity ratio for alternate distributions q_i^k . As the number of records N increases, the allowable difference in entropy of alternative distributions q_i^k varies as $1/N$, given the fixed multiplicity ratio. This is in fact the central issue pointed out in Jaynes (1983): as size of the system increases, the maximum difference in entropy that allowable alternative distributions can have falls like $1/N$, and not $1/\sqrt{N}$ as one might expect.

This forms the basis of a statistical test, outlined in more detail in Jaynes (1983), which can be used for inference purposes. Based on (14) Jaynes names ΔH the entropy difference between the maximum entropy distribution and an alternative distribution, and shows that $2N \times \Delta H$ follows a chi-squared distribution. Given κ the number of constraints imposed on the distribution and $1 - F$ the desired upper tail area, the acceptable entropy interval given the size of the system and the desired level of significance is given by:

$$\Delta H = \frac{\chi_{N-\kappa-1}^2 (1 - F)}{2N} \quad (15)$$

The effect of this is that when N is large, empirical frequencies that differ from the maximum entropy distribution are vanishingly unlikely. Therefore, given knowledge of a certain number of constraints, the only reasonable distribution that one can predict is indeed the one that maximises the information entropy.

3.2 Equilibrium properties of the prediction

The more important issue is whether the methodology generates predictions that are consistent with economic theory. Indeed, Google would expect that the agents are gaining from the transactions carried out on the markets, and would also expect the data in the missed updates to reflect this. In other words, the arrival of a dataset update at time t might signify that for at least one of the N agents there exists an objective, but unobserved, function of the endowment set $f(\{\varepsilon^k\}, t)$ such that:

$$f(\{\varepsilon^k\}, t) \geq f(\{\varepsilon^k\}, t - 1)$$

In the absence of detailed data on the endowments ε_i^k , the maximised entropy function provides a measure of the information that Google could retrieve about this underlying objective process if the communication channel were still available, for example by using the revealed preference algorithms presented in Varian (2006) on the transaction dataset.⁸ As such, if the underlying objective process is a form of utility maximisation, then the maximised entropy function represents the best measure of a welfare function that one can obtain under uncertainty.

Similarly, assuming an underlying objective process provides economic significance to the β^k parameter. As pointed out in Foley (1994), β^k is best understood in terms of a derivative of the entropy function $H(p_i^k)$. This is the standard interpretation of Lagrangian multipliers, but can be derived by totally differentiating $H(p_i^k)$ in (5) and reintroducing the first order condition (11) to substitute for p_i^k :

$$dH(p_i^k) = \alpha^k \sum_{i=1}^{\omega^k} dp_i^k + \beta^k \sum_{i=1}^{\omega^k} \varepsilon_i^k dp_i^k \quad (16)$$

The first term of the equation is equal to zero, because (8) guarantees that the probability distribution always sums to one, therefore the sum of the shifts in probability must equal zero. The second term can be identified by totally differentiating (9). Keeping in mind that the definition of the endowment levels ε_i^k in section 2.1 imply that they are fixed, this gives:

$$d\omega^k = N \sum_{i=1}^{\omega^k} dp_i^k \varepsilon_i^k \quad (17)$$

Combining (16) and (17), one obtains a definition of β^k as the derivative of the entropy function H^k . If H^k provides the information-theoretical measure of an objective welfare function, then β^k measures the gradient of this function with respect to the quantities available, which is consistent with classical theory.

$$N \frac{dH^k}{d\omega^k} = \frac{dS^k}{d\omega^k} = \beta^k \quad (18)$$

⁸It is interesting to note, with respect to the setting mentioned in the introduction, that Professor Varian has been consulting for Google since 2002 and became the chief economist of Google in 2007.

As a result, whenever entropy provides a measure of welfare, then the β^k parameters can be interpreted as marginal welfare functions, or following Foley (1994), as prices. Indeed, relation between the level of endowment ω^k and the β^k parameters already outlined in section 2.3 has the characteristics of a price. It is infinite when the endowment is zero, and falls to zero with limitless availability.

More importantly, trade between two systems equalises the β^k variables, which one would also expect from prices. Let us assume that Google is told that in addition to the first economy, now denoted by subscript A, there exists a second exchange economy, indicated by a subscript B. The aggregate population is $N = N_A + N_B$ and the aggregate endowments are $\omega^k = \omega_A^k + \omega_B^k$. We also define the following probabilities, which are the probability of finding a given endowment level in A and B, as well as the probability that a randomly chosen record belongs to dataset A or B:

$$\left\{ \begin{array}{l} p_{A,i}^k = \frac{n_{A,i}^k}{N_A} \\ p_{B,i}^k = \frac{n_{B,i}^k}{N_B} \end{array} \right. \quad \text{and} \quad \left\{ \begin{array}{l} \rho_A = \frac{N_A}{N} \\ \rho_B = \frac{N_B}{N} \end{array} \right.$$

If Google knows that the two economies are separate, then its best guess of the long-run state of each economy is obtained through the maximisation process outlined above. Because the overall endowments ω_A^k, ω_B^k are conserved within A and B, two sets of distributions (12) are obtained, one for each of the K markets in A and B.

If, however, it is known that trade occurs between A and B, then Google cannot apply the conservation constraint (9) to each economy individually. Therefore, the best it can do is to maximise its uncertainty of the joint state of A and B, subject to a joint conservation constraint. The incoming dataset updates Google misses now apply to either datasets A or B. Therefore its uncertainty as to the content of the missed updates is not just related to which market the updated entry belongs to and which endowment level the entry shifts to, but also to which dataset is updated. The joint entropy of a record is now:

$$H = \rho_A H_A + \rho_B H_B \tag{19}$$

Where the entropy of a record in A and B is given by:

$$\begin{cases} H_A = - \sum_{k=1}^K \sum_{i=1}^{\omega^k} p_{A,i}^k \ln p_{A,i}^k \\ H_B = - \sum_{k=1}^K \sum_{i=1}^{\omega^k} p_{B,i}^k \ln p_{B,i}^k \end{cases} \quad (20)$$

The joint conservation constraint now states that the expected endowment level over both economies A and B is equal to the overall per capita endowment.

$$\sum_{i=1}^{\omega^k} (\rho_A p_{A,i}^k + \rho_B p_{B,i}^k) \varepsilon_i^k = \frac{\omega^k}{N} \quad (21)$$

The joint lagrangian and the first order conditions are given in appendix A.2, but following the procedure in section 2.3, one obtains the following result, which is identical to the initial result of equation (12):

$$p_{A,i}^k = p_{B,i}^k = \frac{e^{-\beta^k \varepsilon_i^k}}{\sum_{i=1}^{\omega^k} e^{-\beta^k \varepsilon_i^k}}$$

One can see that the α_A^k and α_B^k control for the different populations, and that the predicted probability distribution is the same in both economies, with a shared β^k parameter. Furthermore, let us assume that initially $\omega_A^k/N_A > \omega_B^k/N_B$, so that $\beta_A^k < \beta_B^k$. The overall per-capita endowment is a weighted average of the initial endowments:

$$\frac{\omega^k}{N} = \rho_A \frac{\omega_A^k}{N_A} + \rho_B \frac{\omega_B^k}{N_B}$$

Therefore, the following relations hold:

$$\frac{\omega_A^k}{N_A} > \frac{\omega^k}{N} > \frac{\omega_B^k}{N_B} \quad \text{and} \quad \beta_A^k < \beta^k < \beta_B^k$$

From this Google can predict that not only should the β^k variables equalise across the two economies, but also that on each market goods will flow from the low β^k economy to the high β^k economy.

4 Discussion

The discussion of the thought experiment and the suggested methodology covers three main aspects. The first relates to the implications entropy concentration theorem for prediction and inference, while the second is the implication of the information-theoretical methodology with respect to the role of rationality in economics. The final discussion relates to the care that must be taken when applying this methodology, as one can easily reach incorrect conclusions.

4.1 Prediction and inference in the Google thought experiment

The entropy concentration theorem discussed in section 3.1 establishes that the maximum entropy approach provides the best prediction as to the overall equilibrium state of the economy in the absence of a detailed description of that state. This methodology can therefore be used for predictive purposes in a theoretical setting, for example in the context of agent-based modeling, in order to obtain predicted theoretical distributions in complement to existing simulation methods.

However, the entropy concentration theorem also provides a statistical test (15) which allows inference to be carried out. Let us assume that Google eventually restores a partial data connection: the detailed dataset is not available, but the empirical frequencies are transmitted. What can Google conclude if this frequency data does not confirm the predicted distribution, in other words its entropy is beyond the allowable bounds of the test (15)? The answer to this question is provided by Jaynes (1983): any significant difference between the predicted and empirical entropies measures the amount of information that can be extracted from the data in order to improve the knowledge of the constraints on the system, hence improving further the prediction. Again, the setting chosen for the thought experiment is not innocent, as this type of information extraction defines Google's business model.

Beyond the possibility of using prediction and inference to improve the understanding of the constraints on the system, this also justifies the viewpoint argued in the introduc-

tion. The maximum entropy methodology, particularly using the information-theoretical definition of Shannon (1948), is best approached from the subjective point of view of an ignorant observer. The fact that this method provides the observer's sharpest prediction does not imply that it necessarily provides an accurate description of the objective reality, only that it provides the best prediction with respect to existing knowledge about the constraints on the endowment distributions.

4.2 Rationality-free modeling

While this methodology can prove to be useful from the prediction/inference point of view, there is a more fundamental theoretical implication of this thought experiment. Google's best prediction includes a stable long run distribution and a set of market variables that display price-like characteristics. Crucially however, these do not depend on the underlying nature of the updating process described in sections 1 and 2.2. Because this process is approached from the point of view of an ignorant observer, the updates can only be treated as random variables where every possible alternative receives an equal weighting. As outlined above, if the unobserved trading process is non-random and the odds of a transaction occurring depend on the value of the agents' objective functions, then the entropy and β parameters take on an economic interpretation. Their existence, however, does not depend on the nature of the underlying process as even a totally random process still leads to the same prediction.

In this respect the Google thought experiment is a direct analogy to the classical Turing (1950) test, where an observer communicates both with a computer and a human being, and is asked to identify which is which. One could imagine that when Google receives the empirical frequency data through the partial channel, a second set of frequencies is sent which is randomly generated subject to some constraints. The maximum information entropy method tells us is that if the constraints applied in the generation of these random frequencies are identical to those that apply to the observed data, the Google should be unable to distinguish the two. As a result, the information theoretical framework rationalises the contributions of Becker (1962) and Gode and Sunder (1945), which both

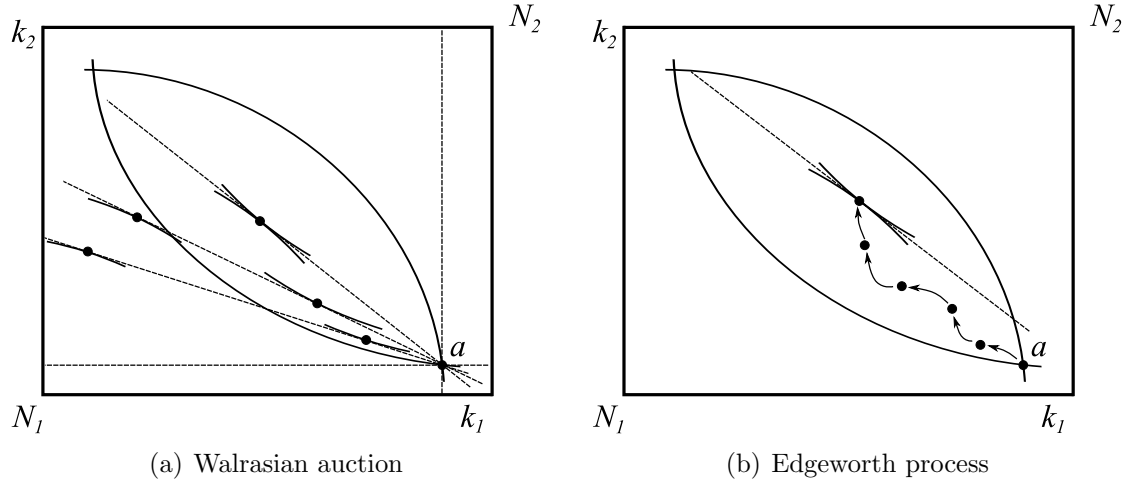


Figure 1: Updating process and endowment conservation

suggest that the efficient operation of markets can be linked to the structure of constraints on agent behaviour rather than the assumed rationality of the behaviour itself.

The implication is that in the absence of detailed information on the state of the economy, knowledge of the determinants of individual agent behaviour, beyond knowledge of the general constraints to the system, cannot improve our prediction. Knowledge of the determinants of agent behaviour is useful in the interpretation of the equilibrium of the system, but it is not needed to predict an equilibrium.

A first setting compatible with this information-theoretical methodology is the Walrasian auction. Here Google's dataset is separate from the auctioneer's and consists of the initial endowments prior to the start of the auctions. The updates that are received by the auctioneer and missed by Google represent the desired endowment level of the agents for each of the quoted prices. Only when the sum of the desired endowment levels satisfies the conservation constraint (9), in other words the excess demand has been eliminated, does the economy reach a Pareto-efficient equilibrium and the auction stops. A practical example of such a setup is provided by the analysis of the Tokyo grain market by Eaves et al. (2008).

An important aspect is that only the initial and equilibrium allocations are within the state space allowed by the endowment conservation constraint (9), which illustrated with the edgeworth box framework of Figure 1. The initial allocation is given by a

and the relative prices quoted by the auctioneer correspond to the rays passing through this point. These are used by the auctioneer to ‘scan’ the endowment space for the equilibrium solution. Whenever the desired endowment levels do not coincide there is excess demand, which violates (9). Therefore, all the intermediate data describing desired allocations as part of the updating process are outside of the allowable state space. This has important implications with regards to how the availability of information determines the conservation of endowments.

A second set of compatible settings are Edgeworth processes and congestion/potential games.⁹ Both Kandori et al. (2008) and Blume (1997) present a market allocation problem similar to the one analysed here and the resulting equilibrium distribution is similar to (12). However, although the problem of allocation is approached from an objective viewpoint, where the behaviour of agents is known (as is the case in Foley (1994)), while the thought experiment rests on a subjective approach which quantifies the ignorance of an observer and makes no assumption on agent behaviour.

In both the analysis of Blume (1997) and Kandori et al. (2008) the equilibrium distribution is obtained using the detailed balance condition of the objective Markov process. This provides the intuition explaining why the predictions of these two studies is compatible with the one made above: the assumption of reversibility required for the existence of a detailed balance condition in a Markov process is the same as the principle of insufficient reason that underpins the use of the entropy concentration theorem by Jaynes (1957a,b). In particular Blume (1997) modifies the detailed balance condition and obtains the following expression for the odds ratio of two states m and n :

$$\frac{\rho^\epsilon(m)}{\rho^\epsilon(n)} = C_{nm} e^{\frac{1}{\epsilon}(P(m)-P(n))} \quad (22)$$

This is directly comparable to the multiplicity ratio (14): $P(m)-P(n)$ is the difference in potential between the two states and ϵ controls the amount of noise in the players

⁹These are discussed together here, because as pointed out by Kandori et al. (2008), their work on Edgeworth processes shares a direct link with the potential game examined in Blume (1997). Furthermore, Monderer and Shapely (1996) show that all congestion games are isomorphic to a potential game, where a single potential function encodes all the payoffs of the agents.

response, in other words the divergence between the best response and the response effectively made. C_{nm} is a constant that depends only on the states m and n . As $\epsilon \rightarrow 0$, the agents always choose their best response, i.e. they do not make mistakes.¹⁰

Allowable deviations from the maximum potential in (22) are controlled in the same way as deviations from maximum entropy in the concentration theorem. This provides an answer to the question raised by Monderer and Shapely (1996) regarding the economic interpretation of the ‘potential’ : if the objective behaviour of agents in the economy corresponds to a congestion game setting, then the information entropy is a measure of the potential and again measures the amount of information one could retrieve about the various payoffs in the system.¹¹

Furthermore the findings of these studies are consistent with the argument that the specification of the equilibrium distribution does not depend on the detailed behaviour of agents. Blume (1997) shows that the stable distribution is the same regardless of the connection of the graph, in other words even when only localised play between neighbouring agents is allowed. Kandori et al. (2008) also show, using the Edgeworth process rather than a potential game, that this result holds in a much more general setting, where each agent has a specific precision parameter, and several agents can change their simultaneously endowments through coalitions. Furthermore, they find that the stable long run distribution, similar to that of Blume (1997) is the same regardless of the size of the noise component in the decision process.

4.3 Avoiding some pitfalls

The first pitfall that must be avoided when using the information theoretical methodology suggested here is to automatically associate increases in entropy with increases in utility.

¹⁰The relation between $1/\epsilon$ in (22) and N in (14) is as follows: in the maximum entropy methodology, the precision is controlled directly by the number of distinct records N . Because the entropy function (5) measures the uncertainty of the observer with respect to the information content a single record, as N increases, the odds of picking a record that deviates from the dominant distribution fall. In Blume’s approach, however, the population is constant and the deviation from the best response is directly controlled through ϵ .

¹¹In particular the entropy function H is the same for all N records, which corresponds to the definition of Monderer and Shapely (1996), who show that the potential function is unique and has to be the same for all players of the game

For example, increases in information entropy due to the updating process occur regardless of the direction of time. If Google’s dataset only describes the economy at time t , then it is as uncertain about the state of the economy n steps into the past as it is n steps into the future.¹² Only if agents’ transactions depend on non-random stepwise improvements is it meaningful to associate increases in entropy to increases in welfare for the relevant direction in time. Therefore, it is best to see the maximum entropy function simply as a measure of the welfare function in an equilibrium state.

The second pitfall, at the other extreme, that the rationality or behaviour of agents does matter at all for market outcomes. In fact, as explained above, the information entropy measure can only be associated with an objective interpretation if the updating process is determined by a systematic but unobserved regularity in the transaction patterns of agents. The maximised entropy is then a measure of the amount of information one would have about this objective function should the detailed state of the system be revealed to an observer. As mentioned above, this type of information extraction underpins Google’s business model, and is directly linked to its privileged access to data compared to the average economic agent.

A final pitfall would be to conclude from the dominance of the maximum entropy prediction that perfectness or imperfectness of information does not matter for the equilibrium. Drawing from the vast literature on this aspect we would expect this imperfectness of information to affect the outcome of market processes. However, the maximum entropy methodology outlined through the thought experiment could suggest that the economy ends up in the same overall state regardless of the underlying nature of the updating process.

Although this may be true from the point of view of the equilibrium frequency of endowments, this is not the case for the detailed equilibrium state. Even in the simple exchange economy framework developed here, the amount of information that individual agents can access matters for the conservation of the value of endowments. This has been addressed by Smith and Foley (2008) using a comparison of the Walrasian and Edgeworth

¹²This is in particular the reason why the reversibility condition mentioned in the previous section, required for a detailed balance condition to hold, is equivalent to adopting a subjective viewpoint.

ontologies similar to the one presented in Figure 1. In the first case, prior to each round of updating the agents receive a quoted price, and in the second no public information is circulated. In the auction the availability information brings a conservation of the real value of endowments, as this value is known to agents prior to effectively trading. In the Edgeworth process, however, prices as a measure of the intrinsic value of goods are formed as a result of the transactions, and given that the initial transactions are made under a state of ignorance there is no guarantee that endowments are conserved. This aspect is relevant if production is to be factored in to this type of analysis, as the existence and capacity of such communication channels will condition the ability of producers to identify production opportunities and therefore coordinate their production decisions.

5 Conclusion

The relevance of the Google thought experiment and the use of the information-theoretical approach is that, as pointed out in the seminal contribution of Hayek (1945), we are all observers whose knowledge of the state of the economy is vanishingly small. Compared to the thought experiment, not only are we uncertain of the process that modifies the initial dataset, but additionally even the initial state is unknown. Even for the case where some data is accessible, it does not completely describe the state of the economy. In such a context, the information-theoretical methodology of Jaynes (1957a,b) uses the information entropy function as a measure of the ignorance of an observer. As a result, following this methodology provides the best prediction of the state of an economy in the absence of all other knowledge than the constraints that apply to agents.

The central conclusion of the thought experiment is that such an observer should not be surprised to learn that overall regularities exist on the markets of an exchange economy, in particular stable endowment distributions and price-like variables that vary with endowment size and equalise with trade. On the contrary, these are some of the few properties that he can accurately predict. If the observer assumes that agents determine their transactions with respect to objective considerations, then the entropy function doubles up as a measure of the information content about these objective processes that

can be retrieved once a dataset becomes available. Therefore, while economic behaviour clearly matters from the point of view of describing the detailed situation, it is not required to predict the existence of a market equilibrium.

References

- Becker, G.S. (1962), Irrational behaviour and economic theory, *The Journal of Political Economy*, 70: 1–13.
- Blume, L.E. (1997), *The Economic System as an Evolving Complex System II*, chap. Population Games, Addison-Wesley.
- Eaves, J., Melvin, M., Mohapatra, S. (2008), Excess demand and price formation during a walrasian auction, *Journal of Empirical Finance*, 15: 533–548.
- Foley, D.K. (1994), A statistical equilibrium theory of markets, *Journal of Economic Theory*, 62: 321–345.
- Gode, D.K., Sunder, S. (1945), Allocative efficiency of markets with zero-intelligence traders: Markets as a partial substitute for individual rationality, *The Journal of Political Economy*, 35: 519–530.
- Hayek, F. (1945), The use of knowledge in society, *The American Economic Review*, 35: 519–530.
- Jaynes, E.T. (1957a), Information theory and statistical mechanics i, *The Physical Review*, 106: 620–630.
- (1957b), Information theory and statistical mechanics ii, *The Physical Review*, 108: 171–190.
- (1983), *E.T. Jaynes, Papers on Probability, Statistics and Statistical Physics*, chap. Concentration of distributions at entropy maxima, Kluwer Academic Publishers, 317–336.

- Kandori, M., Serrano, R., Volij, O. (2008), Decentralised trade, random utility and the evolution of social welfare, *Journal of Economic Theory*, 140: 328–338.
- Monderer, D., Shapely, L.S. (1996), Potential games, *Games and Economic Behaviour*, 14: 124–143.
- Ross, S. (2002), *A First Course in Probability*, Prentice-Hall, 6 ed.
- Shannon, C.E. (1948), A mathematical theory of communication, *The Bell System Technical Journal*, 27: 379–423.
- Smith, E., Foley, D.K. (2008), Classical thermodynamics and economic general equilibrium theory, *Journal of Economic Dynamics and Control*, 32: 7–65.
- Theil, H. (1967), *Economics and Information Theory*, North-Holland.
- Turing, A.M. (1950), Computing machinery and intelligence, *Mind*, 59: 433–460.
- Varian, H. (2006), *Samuelsonian Economics and the 21st Century*, chap. Revealed Preference, Oxford University Press, 99–115.

A Appendix

A.1 Classical vs information entropy

Replacing the multiplicity (2) into the classical definition of entropy (3) gives:

$$S = \sum_{k=1}^K \left(\ln N! - \sum_{i=1}^{\omega^k} n_i^k \right)$$

Given Stirling's approximation, $\ln N! \approx N \ln N - N$. Replacing above one obtains:

$$S \approx \sum_{k=1}^K \left(N \ln N - N - \sum_{i=1}^{\omega^k} (n_i^k \ln n_i^k - n_i^k) \right)$$

The N and $\sum_{i=1}^{\omega^k} n_i^k$ terms cancel out:

$$S \approx \sum_{k=1}^K \left(N \ln N - \sum_{i=1}^{\omega^k} n_i^k \ln n_i^k \right)$$

Isolating $-N$:

$$S \approx -N \sum_{k=1}^K \left(\sum_{i=1}^{\omega^k} \frac{n_i^k}{N} \ln n_i^k - \ln N \right)$$

$$S \approx -N \sum_{k=1}^K \left(\sum_{i=1}^{\omega^k} \frac{n_i^k}{N} \ln n_i^k - \sum_{i=1}^{\omega^k} \frac{n_i^k}{N} \ln N \right)$$

Reintroducing the definition of p_i^k gives the following relation, which corresponds to the equations (4) and (5).

$$S \approx -N \sum_{k=1}^K \sum_{i=1}^{\omega^k} p_i^k \ln p_i^k \quad (\text{A-1})$$

A.2 Trade and β equalisation

The joint entropy of systems A and B is given by (19) and (20). Integrating constraints (8) and (21) gives the following joint lagrangian:

$$L = -\rho_A \sum_{k=1}^K \sum_{i=1}^{\omega^k} p_{A,i}^k \ln p_{A,i}^k - \rho_B \sum_{k=1}^K \sum_{i=1}^{\omega^k} p_{B,i}^k \ln p_{B,i}^k \\ - \sum_{k=1}^K \left(\alpha_A^k \left(\sum_{i=1}^{\omega^k} p_{A,i}^k - 1 \right) + \alpha_B^k \left(\sum_{i=1}^{\omega^k} p_{B,i}^k - 1 \right) + \beta^k \left(\sum_{i=1}^{\omega^k} (\rho_A p_{A,i}^k + \rho_B p_{B,i}^k) \varepsilon_i^k - \frac{\omega^k}{N} \right) \right)$$

This gives the following first order conditions with respect to $p_{A,i}^k$ and $p_{B,i}^k$:

$$\begin{cases} \frac{\partial L}{\partial p_{A,i}^k} = -\rho_A (\ln p_{A,i}^k + 1) - \alpha_A^k - \beta^k \rho_A \varepsilon_i^k = 0 \\ \frac{\partial L}{\partial p_{B,i}^k} = -\rho_B (\ln p_{B,i}^k + 1) - \alpha_B^k - \beta^k \rho_B \varepsilon_i^k = 0 \end{cases}$$

Rearranging:

$$\begin{cases} \ln p_{A,i}^k = -\left(1 + \frac{\alpha_A^k}{\rho_A}\right) - \beta^k \varepsilon_i^k \\ \ln p_{B,i}^k = -\left(1 + \frac{\alpha_B^k}{\rho_B}\right) - \beta^k \varepsilon_i^k \end{cases}$$

$$\begin{cases} p_{A,i}^k = e^{-\left(1 + \frac{\alpha_A^k}{\rho_A}\right)} e^{-\beta^k \varepsilon_i^k} \\ p_{B,i}^k = e^{-\left(1 + \frac{\alpha_B^k}{\rho_B}\right)} e^{-\beta^k \varepsilon_i^k} \end{cases}$$

Replacing this in the two normalisation constraints (19) , one finds that the value of the α_A^k and α_B^k multipliers simply compensates for the difference in the size of the populations N_A and N_B :

$$e^{-\left(1 + \frac{\alpha_A^k}{\rho_A}\right)} = e^{-\left(1 + \frac{\alpha_B^k}{\rho_B}\right)} = \frac{1}{\sum_{i=1}^{\omega^k} e^{-\beta^k \varepsilon_i^k}}$$

A.3 Entropy concentration

According to Stirling's approximation,

$$x! \approx \sqrt{2\pi x} \left(\frac{x}{e}\right)^x$$

Applying this to the multiplicity ratio (13) gives:

$$\frac{W(p_i^k)}{W(q_i^k)} = \prod_{i=1}^{\omega^k} \left(\frac{\sqrt{2\pi N q_i^k} (N q_i^k)^{N q_i^k} e^{-N q_i^k}}{\sqrt{2\pi N p_i^k} (N p_i^k)^{N p_i^k} e^{-N p_i^k}} \right)$$

Rearranging:

$$\frac{W(p_i^k)}{W(q_i^k)} = \prod_{i=1}^{\omega^k} \left(\frac{\sqrt{q_i^k} (N q_i^k)^{N q_i^k} e^{-N q_i^k}}{\sqrt{p_i^k} (N p_i^k)^{N p_i^k} e^{-N p_i^k}} \right)$$

$$\frac{W(p_i^k)}{W(q_i^k)} = \prod_{i=1}^{\omega^k} \left(\sqrt{\frac{q_i^k}{p_i^k}} \frac{e^{N q_i^k \ln N q_i^k} e^{-N q_i^k}}{e^{N p_i^k \ln N p_i^k} e^{-N p_i^k}} \right)$$

$$\frac{W(p_i^k)}{W(q_i^k)} = \prod_{i=1}^{\omega^k} \left(\sqrt{\frac{q_i^k}{p_i^k}} \frac{e^{N q_i^k \ln q_i^k} e^{N q_i^k \ln N} e^{-N q_i^k}}{e^{N p_i^k \ln N p_i^k} e^{N p_i^k \ln N} e^{-N p_i^k}} \right)$$

$$\frac{W(p_i^k)}{W(q_i^k)} = \sqrt{\prod_{i=1}^{\omega^k} \frac{q_i^k}{p_i^k}} \left(\frac{e^{N \sum_{i=1}^{\omega^k} q_i^k \ln q_i^k} e^{N \ln N} e^{-N}}{e^{N \sum_{i=1}^{\omega^k} p_i^k \ln p_i^k} e^{N \ln N} e^{-N}} \right)$$

Using the definition of the information entropy of a probability distribution (5), one obtains:

$$\frac{W(p_i^k)}{W(q_i^k)} = \sqrt{\prod_{i=1}^{\omega^k} \frac{q_i^k}{p_i^k}} \left(\frac{e^{NH(q_i^k)}}{e^{NH(p_i^k)}} \right)$$

$$\frac{W(p_i^k)}{W(q_i^k)} = \sqrt{\prod_{i=1}^{\omega^k} \frac{q_i^k}{p_i^k}} e^{N(H(p_i^k) - H(q_i^k))}$$