Games and Economic Behavior 91 (2015) 1–13

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Competitive economy as a ranking device over networks ☆

Ye Du [a,1], Ehud Lehrer [b,c,2], Ady Pauzner [d,*,3]

[a] School of Finance, Southwestern University of Finance and Economics, China
[b] School of Mathematical Sciences, Tel Aviv University, Tel Aviv 69978, Israel
[c] INSEAD, Bd. de Constance, 77305 Fontainebleau Cedex, France
[d] The Eitan Berglas School of Economics, Tel Aviv University, Tel Aviv 69978, Israel

ARTICLE INFO

Article history:
Received 26 September 2013
Available online 25 March 2015

JEL classification:
D5
D7
D8

Keywords:
Network
Exchange economy
Competitive prices
Ranking
Economy-based ranking
PageRank
Taxation
Citation Count
Normalized citation count

ABSTRACT

We propose a novel approach to generating a ranking of items in a network (e.g., of web pages connected by links or of articles connected by citations). We transform the network into an exchange economy, and use the resulting competitive equilibrium prices of the network nodes as their ranking. The widely used Google's PageRank comes as a special case when the nodes are represented by Cobb–Douglas utility maximizers. We further use the economic metaphor to combine between the Citation Count and PageRank by imposing a redistributive taxing scheme.

1. Introduction

In a world with abundant information, ranking systems are of utmost importance. Well known examples include Google's PageRank, which helps Internet users to identify web pages that are more likely to interest them, and the citation count,[4]

☆ This paper contains the results obtained independently by two groups and replaces "Ranking Via Arrow–Debreu Equilibrium" by Du. The authors wish to thank Elchanan Ben-Porath, Eddie Dekel, Ignacio Palacios-Huerta, David Schmeidler, Daniel Seidman, Orit Tykocinski, Asher Wolinsky and Tim van Zandt for their valuable comments. We are specially indebted to Dudu Lagziel for his help with the simulations.

* Corresponding author.
E-mail addresses: duye@swufe.edu.cn (Y. Du), lehrer@post.tau.ac.il (E. Lehrer), pauzner@post.tau.ac.il (A. Pauzner).

[1] Part of the work was done when Du was a PhD student in the University of Michigan.

[2] Lehrer's research was supported in part by ISF through Grant #538/11.

[3] Pauzner worked on this project also while visiting the Institute for Advanced Studies at the Hebrew University of Jerusalem, and while he was also teaching at the Interdisciplinary Center, Herzlia. His research was supported in part by the Pinhas Sapir Center for Development.

[4] The citation count, which simply counts the number of citations received by a paper, is used to measure the quality of an individual paper. A close related measurement is the famous Impact Factor, which is used to measure the quality of an academic journal. The Impact Factor shares the same spirit of the citation count as it essentially counts the average number of citations received by papers published on it. In this paper, we focus on the problem of ranking individual papers.

http://dx.doi.org/10.1016/j.geb.2015.03.006
0899-8256/

which helps academic researchers to assess the quality of academic articles. In both cases, as well as in many other contexts, the ranking of items is based solely on the information embodied in the network such as links between web sites or citations between scientific articles.

Whereas the citation count merely counts the number of citations from other articles and does not discriminate by the source of the citation, more sophisticated ranking systems attempt to give more weight to votes from items which are of a higher rank (as ascribed by the system itself). The approach taken by PageRank, for example, is based on the idea of translating the link structure to a Markov process as follows: each web page is viewed as a state, and after the random walk hits a state it moves on randomly to one of the states that the current state gives them a link. The ranking of a web page is defined as the long-run proportion of time that the process spends in a given state. Since this value depends not only on the number of incoming links but also on the proportion of time the system spends on the states that send these links, the induced ranking indeed grants more weight to links from higher-ranked pages.

We propose a different approach to ranking that consists of constructing an economy based on the network of links and deriving the ranks from the competitive equilibrium prices.[5] Our approach employs a neoclassical pure-exchange economy. Each web page is represented by one consumer, who is initially endowed with one unit of a specific good. The market price of his good becomes his budget, which in turn serves to buy other goods. The consumer derives utility from consuming the specific goods provided by exactly those consumers that he or she sends a link to. For example, if page i has links only to pages j and k, then consumer i has utility $u(x^i_j, x^i_k)$ where x^i_j, x^i_k are the quantities that i consumes from goods j and k.

The main idea of this paper is to use the competitive prices of this economy as a ranking system. That is, the ranking of a web page is defined as the price of the corresponding good. In this pure-exchange economy, higher-ranked pages correspond to more expensive goods. Moreover, since the budget of a consumer equals the price of the good he initially owns, the initial owners of highly demanded goods are rich. These owners demand larger quantities of the goods they like, thus pushing their prices higher. Hence, those web pages that are pointed to by highly ranked web pages, are highly ranked themselves.

The specific ranking obtained depends on the modeler's choice of utility functions. We assume throughout that all consumers, although consuming different goods, have the same type of utility function. We first consider the Cobb–Douglas utility function which is perhaps the most widely used in economic modeling. We show that when all consumers have Cobb–Douglas preferences, the resulting vector of competitive equilibrium prices coincides with the PageRank ranking system.

We next show that the citation count cannot be derived from our economy. The reason is that, in this type of economy, the value of a paper is identified with the budget of the corresponding consumer, which in turn is his reviewing power. However, in citation count, the reviewing power of a paper is a fixed constant. This impossibility extends also to the 'Normalized' Citation Count, in which the value of a citation from an article is inversely proportional to the number of citations the article makes.

A ranking system based solely on the network information[6] actually considers each item (e.g., web page, article) both as a reviewer, whose judgment (link, references) determines the rank of others, and as a refereed, whose assessed quality depends on the links it obtained from others. The following economic metaphor can sharpen this distinction: the value of an agent as refereed is the market price of its good. Its power as a reviewer equals its budget. In the formulation we employed so far, these two powers are, by definition, the same, since the budget comes from selling the agent's specific good.

The citation count, that does not make any connection between the value of a paper and its refereeing power seems to lack the desirable property of allocating greater weight to citations originating from more important papers. However, there are many cases in which also PageRank, or in fact any ranking system derived from an exchange economy, fails to produce meaningful results. Consider, for example, a sequence of papers published sequentially, each at a distinct time. As citations can only refer to earlier work, all these papers will have zero value. This happens precisely because the value of an item as a reviewer is equated with its value as a refereed entity. The latest article has no incoming links, implying that its price and budget are 0. In turn, it has 0 demand for the articles it has links to. Thus, the penultimate article has 0 value as well, and so forth. In similar cases the citation count – where the reviewing power of all items are the same and independent of their power as refereed items – may perform better than the PageRank.

One can combine the advantages of these two ranking systems by developing the economic metaphor a bit further. We add to the exchange economy a taxation scheme, thus allowing to disentangle the value of an item as a reviewer and its quality as determined by others. Each consumer pays a proportion, say α, of his income as a tax. The tax revenue is then equally redistributed between all the consumers. When $\alpha = 0$, a consumer's budget equals the price of his specific good, leading us back to the original model. With a 100% tax (i.e., $\alpha = 1$), the budgets of all consumers are equal. As a result, their reviewing power is equal regardless of the prices of their goods. The competitive equilibrium prices in this case could serve

[5] More precisely, under some conditions (shown in the paper) the economy has a competitive equilibrium, while under weaker conditions (also shown) the economy is only guarantied to have a quasi-equilibrium (Debreu, 1962). For brevity we refer in both cases to the vector of prices as "competitive prices".

[6] This is as opposed to information obtained from other sources, such as the quality of the journal in which an article was published, etc.