Short Papers

Network Congestion Games Are Robust to Variable Demand

José Correa[1], Ruben Hoeksma[2], and Marc Schröder[3](✉)

[1] Department of Industrial Engineering, University of Chile, Santiago, Chile
correa@uchile.cl

[2] University of Bremen, Bremen, Germany
hoeksma@uni-bremen.de

[3] Chair of Management Science, RWTH Aachen University, Aachen, Germany
marc.schroeder@oms.rwth-aachen.de

Network congestion games have provided a fertile ground for the algorithmic game theory community. Indeed, many of the pioneering works on bounding the efficiency of equilibria use this framework as their starting point. In recent years, there has been an increased interest in studying randomness in this context though the efforts have been mostly devoted to understanding what happens when link latencies are subject to random shocks. In this paper we consider a different source of randomness, namely on the demand side. We look at the basic non-atomic network congestion game with the additional feature that demand is random. Thereto, we introduce an extension of the classic Wardrop equilibrium to fit with this random demand setting. The first obstacle we have to sort out is the definition of equilibrium, as the classic concept of Wardrop equilibrium needs to be extended to the random demand setting. Interestingly, Wang, Doan, and Chen [3], by considering an equilibrium notion in which flow particles evaluate their expected cost using the full knowledge of the demand distribution, conclude that the price of anarchy of the game can be arbitrarily large. In contrast, our main result is that under a very natural equilibrium notion, in which the basic behavioral assumption is that users evaluate their expected cost according to the demand they experience in the system, the price of anarchy of the game is actually the same as that in the deterministic demand game [1, 2]. This is yet another confirmation of the robustness of the price of anarchy to situations in which even the number of players in the system may be random.

A full version of this paper with all the proofs and context can be found at https://www.dii.uchile.cl/~jcorrea/papers/Conferences/CHS2017.pdf.

References

1. Correa, J.R., Schulz, A.S., Stier Moses, N.E.: Selfish routing in capacitated networks. Math. Oper. Res. **29**(4), 961–976 (2004)
2. Roughgarden, T., Tardos, E.: How bad is selfish routing? J. ACM **49**(2), 236–259 (2002)
3. Wang, C., Doan, X., Chen, B.: Price of anarchy for non-atomic congestion games with stochastic demands. Transp. Res. Part B **70**, 90–111 (2014)

N. R. Devanur and P. Lu (Eds.): WINE 2017, LNCS 10674, p. 397, 2017.
https://doi.org/10.1007/978-3-319-71924-5

The Crowdfunding Game

Extended Abstract

Itai Arieli, Moran Koren(✉), and Rann Smorodinsky

Faculty of Industrial Engineering,
Technion - Israel Institute of Technology, Haifa, Israel
{iarieli,ko,rann}@technion.ac.il

The evolution of the 'sharing economy' has made it possible for the general public to invest in early-stage innovative and economically risky projects and products. These funding schemes, dubbed 'crowdfunding', have been gaining popularity among entrepreneurs and it is reported that crowdfunding for supporting new and innovative products has been overwhelming with over 34 Billion Dollars raised in 2015[1].

In addition to serving as an alternative to venture capital funds as a source for fund raising for nascent stage products, the crowdfunding option also serves as a means to gauge market traction for new products. It is implicitly assumed that a successful crowdfunding campaign suggests a high market demand for the new offering.

From the contributor's perspective, the investment in a crowdfunding campaign has two risky aspects. First, the risk of whether the firm will have enough funds to produce and deliver the product; and second, the quality and value of the product is unknown at the time of the campaign and could possibly be disappointing even if eventually delivered.

In many on-line crowd-funding platforms such as "Kickstarter" and "Indiegogo" a typical campaign format has two critical components. First, it sets a price for the future product and second it sets a threshold. Contributions are collected only if in total they exceed this threshold. Both values are determined by the fund raising firm. This format is designed to mitigate the aforementioned risks. If the threshold is set high enough then contributions are collected only when the company has enough funds on the one hand, and the 'wisdom-of-the-crowd' points to a high valued product.

A crowdfunding game, $\Gamma(N, B, p)$, is a game of incomplete information played among a population of N potential contributors (or players). An unknown state of nature $\omega \in \Omega = \{H, L\}$ is drawn with equal prior probabilities. In state H the common value of the product is 1 and in state L it is -1. Conditional on the realized state ω, a private signal $s_i \in S_i = \{H, L\}$ is drawn independently for every player i. We assume $p = Pr(s_i = \omega|\omega) > 0.5$. Each player i has a

The full version can be found at https://arxiv.org/abs/1710.00319.

Rann Smorodinsky—Research supported by GIF research grant no. I-1419-118.4/2017, ISF grant 2018889, Technion VPR grants, the joint Microsoft-Technion e-Commerce Lab, the Bernard M. Gordon Center for Systems Engineering at the Technion, and the TASP Center at the Technion.

[1] Figures taken from http://crowdexpert.com/crowdfunding-industry-statistics.

N. R. Devanur and P. Lu (Eds.): WINE 2017, LNCS 10674, pp. 398–399, 2017.
https://doi.org/10.1007/978-3-319-71924-5

binary action set, $A_i = \{0, 1\}$, with $a_i = 1$ representing a decision to contribute and $a_i = 0$ represents a decision to opt-out and not to contribute. The utility of every player $i \in N$ is defined as follows

$$u_i(a_i, a_{-i}, \omega) = \begin{cases} 1 & \text{if } a_i = 1 \text{ and } \sum_{j \in N} a_j \geq B \text{ and } \omega = H \\ -1 & \text{if } a_i = 1 \text{ and } \sum_{j \in N} a_j \geq B \text{ and } \omega = L \\ 0 & \text{otherwise} \end{cases} . \quad (1)$$

In words, whenever player i chooses not to buy the product, she receives a utility of 0. If she chooses to buy, then her utility is depends on the total number of contributors. If less than B players contributed then the product is not supplied and the utility is once again zero. If it exceeds B then her utility is determined by the state of nature and equals 1 in state H and -1 in state L.

We propose two performance measures for a crowdfunding campaign:

- The *correctness index* of a game is defined as the probability that the game ends up with a the correct decision. That is, the probability the product is funded when its value is 1 or the probability that the product is rejected when its value is -1. The correctness index measures how well the crowdfunding aggregates the private information from the buyers in order to make sure the firm pursues the product only when it is viable.
- The *market penetration index* is the expected proportion of contributors provided that the product is supplied, i.e., the threshold is surpassed. This number serves as a proxy for success of the campaign as a means to attract further investments.

Our theoretical results provide limits on the success, in both aspects, of **large** crowdfunding games. We state and prove three results:

- We provide a constructive proof for the existence of a symmetric, non-trivial equilibrium and we show it is unique. In every such equilibrium players with a high signal surely contribute while those with a low signal either decline or take a mixed strategy whereby they contribute at a positive probability, strictly less than one.
- In large games, we provide a tight bound on the correctness index which is strictly less than one. Thus, no matter how the campaign goal is set, full information aggregation cannot be guaranteed. We compare this with the efficiency guarantees of majority voting implied by Condorcet Jury Theorem.
- I large games, we provide a bound on the penetration index and we show that by setting the champaign goal optimally the resulting market penetration is higher than the benchmark case where the campaign goal is set to a single buyer ($B = 1$).

Calculations, provided in the paper, demonstrate that the asymptotic results approximately hold for small populations of potential contributors.

The Power of Opaque Products in Pricing

Adam N. Elmachtoub and Michael L. Hamilton(✉)

Columbia University, New York, NY 10027, USA
adam@ieor.columbia.edu, mh3461@columbia.edu

Abstract. We study the power of selling opaque products, i.e., products where a feature (such as color, brand, or time) is hidden from the customer until after purchase. Opaque products have emerged as a powerful vehicle to increase revenue for many online retailers, service providers, and travel agents that offer horizontally differentiated items. Customers who are indifferent about the hidden feature typically opt for an opaque product in exchange for a price discount, while customers with strong preferences typically opt for a traditional item at full price. In the models we consider, all traditional items are sold at a single price alongside opaque products corresponding to every possible subset of items. The price of opaque products of the same size are constrained to be the same for practicality. Alternatively, another common approach to increase revenue is to explicitly charge different prices for the items, which we refer to as discriminatory pricing, as opposed to charging one price for all the items, which we refer to as single pricing. In this work, we benchmark the revenue of opaque selling strategies against optimal discriminatory pricing for lower bounds and optimal single pricing for upper bounds. Conceptually, our opaque selling strategy balances the impartiality of single pricing with the price discrimination capabilities of discriminatory pricing.

We consider two types of customer behavior with respect to opaque products, both of which may occur in various applications. Specifically, a customer is called pessimistic if they believe the opaque product will yield their least desired item, and is called risk-neutral if they believe the opaque product will allocate the items with equal probability. In general, we assume customers are unit-demand and utility-maximizing, with i.i.d. item valuations. We show that when customers are pessimistic, opaque selling always dominates discriminatory pricing under any item valuation distribution. When customers are risk-neutral, opaque selling dominates discriminatory pricing in the case where item valuations take only two values (high or low). We also show that opaque selling with just one opaque product can provide up to and at most twice the revenue from single pricing. The revenue increase from having exponentially many opaque products is also at most a constant factor of the revenue from single pricing.

Keywords: Opaque products · Price discrimination · Item pricing

A full version is available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3025944.

N. R. Devanur and P. Lu (Eds.): WINE 2017, LNCS 10674, p. 400, 2017.
https://doi.org/10.1007/978-3-319-71924-5

Information Aggregation in Overlapping Generations

Mohammad Akbarpour[1], Amin Saberi[2], and Ali Shameli[2](✉)

[1] Graduate School of Business, Stanford University, Stanford, USA
mohamwad@stanford.edu

[2] Department of Management Science and Engineering, Stanford University, Stanford, USA
{saberi,shameli}@stanford.edu

Abstract. We create a model of information aggregation with overlapping generations, where agents arrive continuously, meet others over time, share information about an underlying state, and depart at some stochastic time. We examine under what conditions the society will produce individuals with precise knowledge about the state of the world. We consider two information sharing technologies. Under the full information sharing technology, individuals exchange the information about their point estimates of an underlying state, as well as their sources (or the precision of their signals) and update their beliefs by taking a weighted average. Under the limited information sharing technology, agents only observe the information about the point estimates of those they meet, and update their beliefs by taking a weighted average, where weights can depend on the sequence of meetings, as well as the labels and 'ages' of agents they meet. Our main result shows that, unlike static settings, using linear learning rules without access to the precision information will not guide the population (or even a fraction of its members) to converge to a unique belief, and having access to, and exploiting knowledge of the precision of a source signal are essential for having an informed populace.

We would like to thank Michael Harrison, Matthew Jackson, Paul Milgrom, Jeffery Zwiebel, David Kreps, Darrell Duffie, Omer Tamuz, Ben Brooks, Michael Ostrovsky, Alireza Tahbaz-Salehi, Ilan Lobel, Svetlana Bryzgalova, and several seminar participants for valuable conversations.

N. R. Devanur and P. Lu (Eds.): WINE 2017, LNCS 10674, p. 401, 2017.
https://doi.org/10.1007/978-3-319-71924-5

Networked Markets and Relational Contracts

Matt Elliott[1,2], Benjamin Golub[3], and Matt V. Leduc[4,5](✉)

[1] Cambridge University, Cambridge, UK
[2] Caltech, Pasadena, USA
matthew.l.elliott@gmail.com
[3] Harvard University, Cambridge, USA
[4] Aix-Marseille University, CNRS, EHESS, Centrale Marseille, AMSE, Tassin La Demi-Lune, France
ben.golub@gmail.com
[5] IIASA, Laxenburg, Austria
mattvleduc@gmail.com

Abstract. Empirical studies of commercial relationships between firms reveal that (i) suppliers encounter situations in which they can gain in the short run by acting opportunistically—for example, delivering a lower quality than promised after being paid; and (ii) good conduct is sustained not exclusively by formal contracts but through informal relationships and the expectation of future business. In such relationships, the need to offer each supplier a large enough share of future business to deter cheating limits the number of supply relationships each buyer can sustain. The market thus becomes networked, with trade restricted to durable relationships. We propose and analyze a simple dynamic model to examine the structure of such overlapping relational contracts in equilibrium. Due to exogenous stochastic shocks, suppliers are not always able to make good on their promises even if they wish to, and so links are constantly dissolving and new ones are forming to take their place. This induces a Markov process on networks. We study how the stationary distribution over networks depends on the parameters—most importantly, the value of trade and the probability of shocks. When the rate at which shocks hit increases, as might happen during an economic downturn, maintaining incentive compatibility with suppliers requires promising each more future business and this necessitates maintaining fewer relationships with suppliers. This results in a destruction of social capital, and even if the rate of shocks later returns to its former level, it can take considerable time for social capital to be rebuilt because of search frictions. This creates a novel way for shocks to be persistent. It also suggests new connections between the theory of relational contracting, on the one hand, and the macroeconomic analysis of recessions, on the other.

Paper available at http://ssrn.com/abstract=3049512.

N. R. Devanur and P. Lu (Eds.): WINE 2017, LNCS 10674, p. 402, 2017.
https://doi.org/10.1007/978-3-319-71924-5

On Variants of Network Flow Stability

Young-San Lin[1(✉)] and Thanh Nguyen[2]

[1] Computer Science Department, Purdue University, West Lafayette, USA
lin532@purdue.com

[2] Krannert School of Management, Purdue University, West Lafayette, USA
nguye161@purdue.edu

Abstract. In a stable network flow problem, we are given a directed and capacitated network, where each vertex has strict preference over their incident edges, and need to find a flow between a source and a sink that is stable with respect to deviations along any path. A common interpretation of this problem is that the vertices represent agents and the edges represent potential contracts between the endpoint agents; a directed edge from an agent A to an agent B represents the possibility of agent B buying products via a contract from agent A. A stable flow is an equilibrium trade pattern, where no group of agents can all benefit from rerouting the flow along a path among themselves.

The stable flow problem is well studied and has several applications in supply chain and trading networks. However, the Kirchhoffs law, which requires the inflow is equal to the outflow for every vertex of the network, limits the applicability of this problem. For example, in a supply chain network, one vertex can represent a manufacturing firm that takes raw materials as input and produces certain part-products while another vertex might correspond to an assembly firm whose inputs are the part-products and outputs are finished products. Clearly, the Kirchhoffs law does not hold for both manufacturing and assembly nodes in this example.

In this paper, we consider a generalization of the traditional stable flow problem, in which the outflow is monotone piecewise linear to the inflow for each vertex. We first show the existence of flow stability by reducing this variant of stable flow problem to Scarf's Lemma, then introduce a path augmenting algorithm that runs in polynomial time.

We first define a monotone piecewise linear mapping network (MPLM-network). A convex monotone piecewise linear mapping network (CMPLM-network) is defined as a subcategory of MPLM-networks where the slopes of the piecewise linear functions are in increasing order for every agent. A linear mapping network (LM-network) is a subcategory of CMPLM-networks where the amount of outgoing contracts of every agent with incoming contracts is a linear function on the amount of incoming contracts.

A flow assignment is stable if there does not exist a blocking path in a network. A flow assignment has a blocking path P if the first agent in P prefers to offer contracts to the second agent in P to some other agents she had already offered, while intermediate agents still have space for signing contracts, and the last agent in P prefers to accept the contracts

N. R. Devanur and P. Lu (Eds.): WINE 2017, LNCS 10674, pp. 403–404, 2017.
https://doi.org/10.1007/978-3-319-71924-5

offered by the penultimate agent in P to some other agents she had already accepted. The existence of stable flow in CMPLM-networks can be proved by a reduction to Scarf's Lemma. LM-networks, as a subcategory of CMPLM-networks, always have a stable flow assignment. Every MPLM-network has a corresponding LM-network by transforming each agent into a subnetwork. Therefore, stability always exists in MPLM-networks.

A constructive way to find a stable flow in acyclic LM-networks is similar to the path augmenting algorithm for the original stable flow problem. The approach is a variant of deferred acceptance algorithm among agents. The main difference is in LM-networks, flow conservation no longer holds. As a result, in each path augmenting iteration, we augment along a path from the source agent to the sink agent, or along a σ-cycle, a path from the source agent to a cycle. The running time for LM-network is $O(|V||E|)$. For MPLM-networks, the running time of our algorithm is $O(|V|(|E|+K))$ where K is the total number of piecewise linear segments.

For each cyclic LM-network, there exists an equivalent acyclic LM-network consists of the source vertex, the sink vertex, and three layers of vertices between the source and the sink. Hence, we can constructively find a stable flow in a cyclic LM-network by reducing this instance to its equivalent acyclic LM-network. The numbers of vertices and edges in the acyclic network are just a constant factor of that in the cyclic network. The running time is $O(|V||E|)$ and similarly the running time for cyclic MPLM-networks, the running time is $O(|V|(|E|+K))$.

Acknowledgement. This research is partly supported by National Science Foundation Grants AST- 1443965, CMMI 1728165.
The full paper can be found at: https://arxiv.org/pdf/1710.03091.pdf.

Balancing Efficiency and Equality in Vehicle Licenses Allocation

Zhou Chen[1], Qi Qi[1(✉)], and Changjun Wang[2]

[1] The Hong Kong University of Science and Technology, Clear Water Bay, Hong Kong
{chenaq,kaylaqi}@ust.hk
[2] Beijing University of Technology, Beijing, China
wcj@bjut.edu.cn

With the global increase of urbanization, the population of urban areas is growing rapidly. Concurrently, the number of private vehicles in these places is increasing dramatically. Due to traffic and air quality concerns, many big cities have begun to adopt the vehicle licenses quantitative control policies. In these cities, a limited number of vehicle licenses are allocated among a very large number of potential car buyers every one or two months. Then how to design an effective mechanism to allocate the limited license quotas becomes a challenging problem. The current allocation mechanisms differ from city to city. Several mechanisms have been developed and implemented in reality, such as auction, lottery, lottery with reserved price, and the simultaneous auction and lottery.

In this work, we target to design the optimal mechanism to balance efficiency and equality in practice. We first propose a unified two-group mechanism framework that either includes or outperforms all the existing mechanisms. Besides, the unified framework also leads to easy implementation in reality due to its truthfulness and simple structure. Under this framework, assuming the players' private values are drawn independently from a common distribution, we prove the optimal mechanism is always sequential auction and lottery. Besides, the optimal allocation rule depends only on the total number of players and the total number of licenses for all commonly used distributions. We then extend the two-group framework to a general multi-group framework. The experimental results show us the optimal two-group mechanism is the best choice in practice. Thus, our work provides an effective tool for social planner to design truthful mechanisms to maximize the social efficiency under any equality level. We also discuss possible applications of our result to resource allocation in other settings.

A full version of this paper is available at http://ssrn.com/abstract=3049504.
This work is partly supported by the Research Grant Council of Hong Kong (GRF Project No. 16213115, 16243516 and 16215717), and by the National Natural Science Foundation of China (NSFC-11601022).

N. R. Devanur and P. Lu (Eds.): WINE 2017, LNCS 10674, p. 405, 2017.
https://doi.org/10.1007/978-3-319-71924-5

Author Index