Abstracts

How to Hide in a Network?

Francis Bloch[1], Bhaskar Dutta[2], and Marcin Dziubiński[3(✉)]

[1] Université Paris 1 and Paris School of Economics,
48 Boulevard Jourdan, 75014 Paris, France
francis.bloch@univ-paris1.fr

[2] University of Warwick and Ashoka University, Coventry CV4 7AL, UK
b.dutta@warwick.ac.uk

[3] Institute of Informatics, University of Warsaw, Banacha 2, 02-097 Warsaw, Poland
m.dziubinski@mimuw.edu.pl

We propose a model of strategic hiding in a network in face of a hostile authority. Given a set of nodes, the hider chooses a network over these nodes together with a node. The network chosen by the hider is observed by the seeker (the hostile authority) but the location choice is not observed. The seeker chooses one of the nodes in the network to inspect. The inspected node is removed from the network. If the hider hides in the inspected node or one of its neighbours, he is caught by the seeker and suffers a penalty. Otherwise, he enjoys the benefits from the network that are a convex and increasing function of the number of nodes (including himself) that the hider can access (directly or not) in the network. This form of network benefits, first proposed by [3], is in line with the celebrated Metcalfe's law, where the function is identity. The objectives of the seeker are to minimize the payoff of the hider and the proposed model takes the form of a two-stage zero-sum game.

The hide and seek stage in our model is similar to the hide and seek games on graphs of [2], with the difference that in their case the penalty from being caught is 0 and benefits from not being caught are fixed and independent of the graph. Unlike in the model of [1], in our model the authorities choose their seeking strategy knowing the network and only one node chooses the network topology to hide himself. This is similar to the model of [4]. However, unlike in their model, the authorities are strategic and they take into account the incentives and strategic behaviour of the hider when choosing the seeking strategy. Although very stylised and simple, the model allows us to capture the trade-off between secrecy and network benefits.

We provide optimal networks for the hider and characterize optimal strategies of the two players on these networks. In general, the optimal networks consists of a number of singleton nodes and a connected component which is either a cycle or a core-periphery network. If the component is a cycle, in equilibrium the hider mixes uniformly across its nodes. If the component is a core-periphery network, the hider mixes uniformly across the periphery nodes. This provides theoretical support to the claim that the hider chooses networks where his centrality is small and indistinguishable from the centralities of the other nodes.

This work was supported by Polish National Science Centre through Grant 2014/13/B/ST6/01807.

© Springer Nature Switzerland AG 2018
G. Christodoulou and T. Harks (Eds.): WINE 2018, LNCS 11316, pp. 441–442, 2018.
https://doi.org/10.1007/978-3-030-04612-5

Full version of the paper is available at: http://www.mimuw.edu.pl/~amosild/hnsnet-full.pdf.

References

1. Baccara, M., Bar-Isaac, H.: How to organize crime. Rev. Econ. Stud. **75**(4), 1039–1067 (2008)
2. Fisher, D.: Two person zero-sum games and fractional graph parameters. Congressus Numerantium **85**, 9–14 (1991)
3. Goyal, S., Vigier, A.: Attack, defence, and contagion in networks. Rev. Econ. Stud. **81**(4), 1518–1542 (2014)
4. Waniek, M., Michalak, T., Wooldridge, M., Rahwan, T.: Hiding individuals and communities in a social network. Nat. Hum. Behav. **2**, 139–147 (2018)

Multiplicative Pacing Equilibria in Auction Markets

Vincent Conitzer[1], Christian Kroer[2(✉)], Eric Sodomka[2], and Nicolas E. Stier-Moses[2]

[1] Econorithms LLC, and Duke University, Durham, USA
vincent.conitzer@duke.edu
[2] Core Data Science, Facebook Inc., Cambridge, USA
{chrkroer,sodomka,nstier}@fb.com

Abstract. Budgets play a significant role in real-world sequential auction markets such as those implemented by Internet companies. To maximize the value provided to auction participants, spending is smoothed across auctions so budgets are used for the best opportunities. Motivated by a mechanism used in practice by several companies, this paper considers a smoothing procedure that relies on *pacing multipliers*: on behalf of each bidder, the auction market applies a factor between 0 and 1 that uniformly scales the bids across all auctions. Reinterpreting this process as a game between bidders, we introduce the notion of *pacing equilibrium*, and prove that they are always guaranteed to exist. We demonstrate through examples that a market can have multiple pacing equilibria with large variations in several natural objectives. We show that pacing equilibria refine another popular solution concept, competitive equilibria, and show further connections between the two solution concepts. Although we show that computing either a social-welfare-maximizing or a revenue-maximizing pacing equilibrium is NP-hard, we present a mixed-integer program (MIP) that can be used to find equilibria optimizing several relevant objectives. We use the MIP to provide evidence that: (1) equilibrium multiplicity occurs very rarely across several families of random instances, (2) static MIP solutions can be used to improve the outcomes achieved by a dynamic pacing algorithm with instances based on a real-world auction market, and (3) for our instances, bidders do not have an incentive to misreport bids or budgets provided there are enough participants in the auction.

Keywords: ad auctions · Budget constraints · Internet advertising Pacing · Repeated auctions

This work was done while the first author was visiting Facebook Core Data Science. A full version of this paper can be found at https://arxiv.org/abs/1706.07151.

© Springer Nature Switzerland AG 2018
G. Christodoulou and T. Harks (Eds.): WINE 2018, LNCS 11316, p. 443, 2018.
https://doi.org/10.1007/978-3-030-04612-5

Trading Networks with Bilateral Contracts

Tamás Fleiner[1], Zsuzsanna Jankó[2], Akihisa Tamura[3], and Alexander Teytelboym[4(✉)]

[1] Budapest University of Technology and Economics, Műegyetem rkp. 3, Budapest 1111, Hungary
fleiner@cs.bme.hu

[2] Corvinus University, Fővám tér 8, Budapest 1093, Hungary
jzsuzsy@cs.elte.hu

[3] Department of Mathematics, Keio University, 3-14-1 Hiyoshi, Kohoku-ku, Yokohama 223-8522, Japan
aki-tamura@math.keio.ac.jp

[4] Department of Economics, University of Oxford, 10 Manor Road, Oxford OX1 3UQ, UK
alexander.teytelboym@economics.ox.ac.uk

Abstract. We consider a model of matching in trading networks in which firms can enter into bilateral contracts. In trading networks, *stable* outcomes, which are immune to deviations of arbitrary sets of firms, may not exist. We define a new solution concept called *trail stability*. Trail-stable outcomes are immune to consecutive, pairwise deviations between linked firms. We show that any trading network with bilateral contracts has a trail-stable outcome whenever firms' choice functions satisfy the full substitutability condition. For trail-stable outcomes, we prove results on the lattice structure, the rural hospitals theorem, strategy-proofness, and comparative statics of firm entry and exit. We also introduce *weak* trail stability which is implied by trail stability under full substitutability. We describe relationships between the solution concepts.

Keywords: Matching markets · Market design · Trading networks
Supply chains · Trail stability · Weak trail stability · Chain stability
Stability · Contracts

Fleiner: Research was supported by the OTKA K108383 research project and the MTA-ELTE Egerváry Research Group; part of the research was carried out during two working visits at Keio University. Jankó: Research was supported by the OTKA K109240 research project and the MTA-ELTE Egerváry Research Group. Tamura: Research was supported by Grants-in-Aid for Scientific Research (B) from JSPS. Teytelboym: This work was supported by the Economic and Social Research Council grant number ES/R007470/1.

© Springer Nature Switzerland AG 2018
G. Christodoulou and T. Harks (Eds.): WINE 2018, LNCS 11316, p. 444, 2018.
https://doi.org/10.1007/978-3-030-04612-5

Markets for Public Decision-Making

Nikhil Garg, Ashish Goel, and Benjamin Plaut(✉)

Stanford University, Stanford, CA, USA
{nkgarg,ashishg,bplaut}@stanford.edu

A public decision-making problem consists of a set of issues, each with multiple possible alternatives, and a set of competing agents, each with a preferred alternative for each issue. We study adaptations of market economies to this setting, focusing on binary issues. Issues have prices, and each agent is endowed with artificial currency that she can use to purchase probability for her preferred alternatives (we allow randomized outcomes). We first show that when each issue has a single price that is common to all agents, market equilibria can be arbitrarily bad.

This negative result motivates a different approach. We present a novel technique called *pairwise issue expansion*, which transforms any public decision-making instance into an equivalent Fisher market, the simplest type of private goods market. This is done by expanding each issue into many goods: one for each pair of agents who disagree on that issue. We show that the equilibrium prices in the constructed Fisher market yield a *pairwise pricing equilibrium* in the original public decision-making problem which maximizes Nash welfare. More broadly, pairwise issue expansion uncovers a powerful connection between the public decision-making and private goods settings; this immediately yields several interesting results about public decisions markets, and furthers the hope that we will be able to find a simple iterative voting protocol that leads to near-optimum decisions.

The full version of the paper can be found at https://arxiv.org/pdf/1807.10836.pdf.

Acknowledgements. This research was supported in part by NSF grant CCF-1637418, ONR grant N00014-15-1-2786, and the NSF Graduate Research Fellowship under grants DGE-114747 and DGE-1656518.

© Springer Nature Switzerland AG 2018
G. Christodoulou and T. Harks (Eds.): WINE 2018, LNCS 11316, p. 445, 2018.
https://doi.org/10.1007/978-3-030-04612-5

Information Signal Design for Incentivizing Team Formation (Extended Abstract)

Chamsi Hssaine(✉) and Siddhartha Banerjee

School of Operations Research and Information Engineering,
Cornell University, Ithaca, NY 14853, USA
{ch822,sbanerjee}@cornell.edu

Abstract. We study the use of Bayesian persuasion (i.e., strategic use of information disclosure/signaling) in endogenous team formation. This is an important consideration in settings such as crowdsourcing competitions, open science challenges and group-based assignments, where a large number of agents organize themselves into small teams which then compete against each other. A central tension here is between the strategic interests of agents who want to have the highest-performing team, and that of the principal who wants teams to be balanced. Moreover, although the principal cannot choose the teams or modify rewards, she often has additional knowledge of agents' abilities, and can leverage this information asymmetry to provide signals that influence team formation. Our work uncovers the critical role of self-awareness (i.e., knowledge of one's own abilities) for the design of such mechanisms. For settings with two-member teams and binary-valued agents, we provide signaling mechanisms which are asymptotically optimal when agents are agnostic of their own abilities. On the other hand, when agents are self-aware, then we show that there is no signaling mechanism that can do better than not releasing information, while satisfying agent participation constraints.

Our work focuses on the use of strategic signaling for incentivizing team formation. The main idea is that many strategic settings have an inherent information asymmetry, where the principal has more information than the participating agents. We seek to understand if there is any way of leveraging this information asymmetry to *influence endogenous team formation*, with the objective of creating balanced teams.

We consider a setting with n agents who form teams of two, leading to some utility for both the agents and the principal. The teams are chosen endogenously by the agents, in the form of a *stable matching*; the principal however can influence agents' preferences via strategic release of information. Each agent has an intrinsic (numerical) type, drawn from some publicly-known prior. Crucially, we assume that each agent's type is *known to the principal*, but *unknown to other agents*. Moreover, an agent's utility is an increasing function of her and her teammates' types, while the principal's utility function depends on the set of resulting teams, and favors having more 'balanced' teams; thus, the principal's and agents'

© Springer Nature Switzerland AG 2018
G. Christodoulou and T. Harks (Eds.): WINE 2018, LNCS 11316, pp. 446–447, 2018.
https://doi.org/10.1007/978-3-030-04612-5

incentives are misaligned. We focus on settings with a binary type-space $\{0,1\}$, and a constant number (K) of prior distributions.

Any signaling policy designed by the principal must induce a *stable matching* of agents, as well as obey *individual rationality* constraints, which enforce that each agent be weakly better off by agreeing to receive the signal. We show that it is enough to restrict to signals that are rank-orderings of agents according to expected posterior types. For $K \geq 1$ prior distributions, we propose the *Cluster First Best* signaling policy, in which agents with types drawn from the same prior distribution are always matched to agents of opposing realized type (i.e., high-type agents are always matched to low-type agents, and vice versa). Our main results are the following:

Theorem 1. *When agents* do not *know their own types, Cluster First Best is asymptotically optimal in n.*

Theorem 2. *When agents* do *know their own types, no signaling policy can do better than random matching.*

Our results indicate the importance of self-awareness in determining the success of signaling mechanisms. Showing this strategy is asymptotically optimal requires a novel dual-certification argument, which may be useful in related settings. Moreover, our work provides important insights and techniques for the design of Bayesian persuasion schemes for general team formation settings, as well as more general bipartite matching settings. For details, refer to our full version: https://arxiv.org/abs/1809.00751.

Acknowledgment. We acknowledge support from the National Science Foundation (grant DMS 1839346), and the Army Research Office (grant W911NF-17-1-0094).

Revenue Management on an On-Demand Service Platform

Vijay Kamble(✉)

Department of Information and Decision Sciences College of Business Administration, University of Illinois at Chicago, Chicago, USA
kamble@uic.edu

Abstract. I consider the optimal hourly (or per-unit-time in general) pricing problem faced by a worker (or a service provider) on an on-demand service platform. Service requests arriving while the worker is busy are lost forever. Thus, the optimal hourly prices need to capture the average hourly opportunity costs incurred by accepting jobs. Due to potential asymmetries in these costs, price discrimination across jobs based on duration, characteristics of the arrival process, etc., may be necessary for optimality, even if the customers' hourly willingness to pay is believed to be identically distributed. I first establish that such price discrimination is not necessary if the customer arrival process is Poisson: in this case, the optimal policy charges an identical hourly rate for all jobs. This result holds even if the earnings are discounted over time. I then consider the case where the customers belong to different classes that are differentiated in their willingness to pay. I present a simple and practical iterative procedure to compute the optimal prices in this case under standard regularity assumptions on the distributions of customer valuations.

Keywords: Optimal pricing · On-demand services

A full draft of the paper is available at https://arxiv.org/abs/1803.06797.

© Springer Nature Switzerland AG 2018
G. Christodoulou and T. Harks (Eds.): WINE 2018, LNCS 11316, p. 448, 2018.
https://doi.org/10.1007/978-3-030-04612-5

Implementing the Lexicographic Maxmin Bargaining Solution

Ashish Goel and Anilesh K. Krishnaswamy(✉)

Stanford University, Stanford, USA
{ashishg,anilesh}@stanford.edu

Abstract. A major question which bargaining theory deals with is that of implementation – designing a mechanism for which a desired bargaining solution is the unique subgame perfect outcome, with each player having complete information, and the social planner/designer having no knowledge of the players' preferences. There has been much work on exhibiting mechanisms that implement various bargaining solutions, in particular the Kalai-Smorodinsky solution and the Nash Bargaining solution. However, to the best of our knowledge, there is no known (subgame perfect) implementation of the lexicographic maxmin solution.

The lexicographic maxmin solution is obtained by a repeated application of the maxmin criterion: first, selecting feasible outcomes that maximize the utility of the worst-off player, then, among these outcomes, selecting those that maximize the utility of the next worst-off player, and so on. The utility gains are measured with respect to the disagreement point. The lexicographic maxmin solution has also had a long history outside of the literature on bargaining. It corresponds directly to the notion of maxmin fairness which has been extensively studied in network routing, bandwidth allocation and other resource allocation problems.

This paper is devoted to designing a mechanism for the (subgame perfect) implementation of the lexicographic maxmin solution. We do so by first defining the *Knockout* mechanism on any two given outcomes. This construction is based on a novel notion, namely *disagreement dominance* (a relation defined on pairs of vectors), which we believe is interesting in its own right. We then use the Knockout mechanism as a subroutine in constructing our full mechanism: a binary tree of games, where each node corresponds to a Knockout mechanism with two outcomes. The workings of our overall mechanism rely crucially on an original combinatorial result we establish, that the lexicographic maxmin solution disagreement dominates any other outcome.

A full draft of the paper is available at https://arxiv.org/pdf/1810.01042v1.pdf.

Ashish Goel—Supported in part by NSF grant CCF-1637418, ONR grant N00014-15-1-2786, and ARO grant W911NF-14-1-0526.
Anilesh K. Krishnaswamy—Also supported by the grants mentioned above.

© Springer Nature Switzerland AG 2018
G. Christodoulou and T. Harks (Eds.): WINE 2018, LNCS 11316, pp. 449–450, 2018.
https://doi.org/10.1007/978-3-030-04612-5

Our mechanism uses the standard assumption that the space of outcomes is such that in any player's best outcome, all the surplus goes to her, and every one else gets no utility. This assumption is commonplace in most of the literature on implementation of bargaining solutions.

Keywords: Bargaining · Implementation · Maxmin fairness
Mechanism design

Exploration vs. Exploitation in Team Formation

Ramesh Johari[1], Vijay Kamble[3(✉)], Anilesh K. Krishnaswamy[2], and Hannah Li[1]

[1] Department of Management Science and Engineering, Stanford University, Stanford, USA
{rjohari,hannahli}@stanford.edu

[2] Department of Electrical Engineering, Stanford University, Stanford, USA
anilesh@stanford.edu

[3] Department of Information and Decision Sciences, University of Illinois at Chicago, Chicago, USA
kamble@uic.edu

Abstract. Modern labor platforms face the online learning problem of optimizing matches between jobs and workers of unknown abilities. This problem is complicated by the rise of complex jobs on these platforms that require teamwork, such as web development and product design. Successful completion of such a job depends on the abilities of all workers involved, which can only be indirectly inferred by observing the aggregate performance of the team. Observations of the performance of various overlapping teams induce correlations between the unknown abilities of different workers at any given time. Tracking the evolution of this correlation structure across a large number of workers on the platform as new observations become available, and using this information to adaptively optimize future matches, is a challenging problem.

To study this problem, we develop a stylized model in which teams are of size 2 and each worker is drawn i.i.d. from a binary (good or bad) type distribution. Under this model, we analyze two natural settings: when the performance of a team is dictated by its strongest member and when it is dictated by its weakest member. We find that these two settings exhibit stark differences in the trade-offs between exploration (i.e., learning the performance of untested teams) and exploitation (i.e., repeating previously tested teams that resulted in a good performance). We establish fundamental regret bounds and design near-optimal algorithms that uncover several insights into these tradeoffs.

Keywords: Team formation · Online learning · Online labor platforms

---

A full draft of the paper is available at https://arxiv.org/abs/1809.06937.

R. Johari—Supported by NSF Grant No. CNS-1544548 and the Stanford Cyber Initiative.

V. Kamble—Supported in part by ARO grant W911NF-14-1-0526.

A. K. Krishnaswamy—Supported in part by NSF grant CCF-1637418, ONR grant N00014-15- 1-2786, and ARO grant W911NF-14-1-0526.

H. Li—Supported by NSF Grant No. CNS-1544548, the Stanford Cyber Initiative, and the Stanford Management Science & Engineering Graduate Fellowship.

© Springer Nature Switzerland AG 2018

G. Christodoulou and T. Harks (Eds.): WINE 2018, LNCS 11316, p. 451, 2018.
https://doi.org/10.1007/978-3-030-04612-5

Author Index