

# Leading Dynamics to Good Behavior

MARIA FLORINA BALCAN  
Georgia Institute of Technology

---

In this note we report recent results on the problem of leading natural dynamics to good behavior in games that have both high-quality and low-quality equilibria. We show how a central agency can use a public-service advertising campaign to help “nudge” players’ behavior towards a high-quality equilibrium, even if only a fraction of players pay attention. We also discuss results analyzing how well-motivated learning rules, when given additional global information about a game, can be used to effectively reach high-quality equilibria.

Categories and Subject Descriptors: F.2 [Analysis of Algorithms and Problem Complexity]: General

General Terms: Algorithms, Theory, Economics

Additional Key Words and Phrases: game theory, dynamics, learning, equilibrium quality

---

## 1. INTRODUCTION

One major focus of algorithmic game theory has been analyzing the *quality* of Nash equilibria in a game, in order to understand the additional cost incurred when we assume that agents are behaving in a strategic manner, rather than performing a global optimization. This focus was initiated by Koutsoupias and Papadimitriou [Koutsoupias and Papadimitriou 1999] who proposed the notion of *Price of Anarchy (PoA)* of a game, which is the ratio of the cost of the worst Nash equilibrium to the social optimum [Nisan et al. 2007]. The PoA has since been studied for a large variety of games, including routing, network design with cost-sharing, job scheduling, network creation etc. (see [Nisan et al. 2007]). While the PoA considers the worst equilibrium, the *Price of Stability (PoS)* [Anshelevich et al. 2004; Correa et al. 2004] instead considers the ratio of the *best* Nash equilibrium to the social optimum. In fact, for many natural games the Price of Anarchy may be very large while the Price of Stability is quite low. For example, in job scheduling on unrelated machines the Price of Anarchy is unbounded while there always is a Nash equilibrium that is also socially optimal (and hence the Price of Stability is 1). Another example is fair cost-sharing games with  $n$  players, where the Price of Anarchy is  $\Theta(n)$  and yet the Price of Stability is only  $\Theta(\log n)$ .

In cases where there are both high-cost and low-cost Nash equilibria, a central authority (such as a government agency) could hope to “nudge” behavior that is currently stuck at a high cost equilibrium into a low cost one by running a public service advertising campaign promoting the better behavior. For example, consider the case of consensus games. Here, players make individual decisions (say of which of two standards to follow) and incur a cost for each neighbor whose decision disagrees with their own. Such a game can have a very high-cost equilibrium where

---

Author’s address: ninamf@cc.gatech.edu

different groups make different decisions and yet have no incentive to change because each player has more neighbors in their group than outside. In this case a central authority might hope to pick one (arbitrary) single decision and advertise that as the desired choice. Indeed, if *everyone* then followed the given advice, then we would immediately achieve a low-cost equilibrium.

However, it is a bit much to assume that everyone will follow any given piece of advice, even if the behavior is optimal if everyone else follows it as well. Instead, a more realistic assumption is that we might advertise a certain behavior, and then only some fraction  $0 < \alpha < 1$  of players will actually go along with it, with the others behaving selfishly. In recent work [Balcan et al. 2009] we consider the question of what can be done in such a scenario, where a public authority would like to encourage behavior to move from a bad equilibrium to a good one, but can only count on some fraction of players acting in the desired manner and who will do so only temporarily. The main question we address in this model is: Is affecting a small constant fraction  $\alpha$  of players sufficient to cause the rest to head towards a low-cost equilibrium, or on the other hand, is even a small constant fraction *not* paying attention enough to cause the whole thing to unravel? In [Balcan et al. 2010], we extend this to the case where rather than a fixed random fraction of players following the given advice, the set of players following it changes over time as players individually test out the advice and use adaptive learning behavior to determine what is individually best for them.

## 2. FORMAL MODELS AND OUR RESULTS

In our basic *advertisement model* [Balcan et al. 2009], the authority first suggests to each player a proposed action, and each player accepts the proposal with some (constant) probability  $\alpha$ . The players that accept the new action are called the *receptive* players (since they are receptive to the advertising campaign), and they stay with the new action long enough so that the *non-receptive* players converge to some Nash equilibrium in the game conditioned on the fixed behavior of the receptive players.<sup>1</sup> Then all players follow an arbitrary best response dynamics from that point on. The goal of the authority is to design its advertising such that this process results in a low-cost equilibrium<sup>2</sup> (we only consider games where best response dynamics is guaranteed to converge to a pure Nash equilibrium).

What we show is that interestingly a number of well-studied games fall into three different categories with respect to the ability of the authority to affect the quality of behavior. Some behave quite nicely, with the authority able to bring the system to a low-cost equilibrium for any constant  $\alpha > 0$ ; some have a threshold property (e.g., producing good equilibria for  $\alpha > 1/2$  but poor equilibria for  $\alpha < 1/2$ ) and some perform poorly for all constants  $\alpha < 1$ . Specifically, we show the following:

- For fair cost sharing games [Anshelevich et al. 2004], if the authority advertises the socially optimal solution, then for any  $\alpha$  the expected cost of the final

<sup>1</sup>We also examine a variation in which we replace the assumption that non-receptive players settle to a Nash equilibrium with their instead performing a small number of rounds of best-response dynamics.

<sup>2</sup>The goal is to get to a low-cost equilibrium irrespective of the quality of the original starting point of our dynamics.

equilibrium is at most an  $O((\log n)/\alpha)$  factor larger than optimal. Recall that the price of stability is  $\Theta(\log n)$  for this game, so the difference in guarantees is only  $O(1/\alpha)$ , whereas the Price of Anarchy is  $\Theta(n)$ . These results extend to the case where we add to the cost model a linear latency function which depends on the load observed on the edges. The results extend also to the case that the authority advertises an approximation to the optimal solution, with a proportional degradation of the final guarantee.

- For job scheduling on unrelated machines [Nisan et al. 2007; Even-Dar et al. 2007] we present a negative result, showing that for  $m = n$  machines, even if the authority can pick  $n - 2$  of the  $n$  jobs to follow the advertising, there is still a possibility of reaching a pure Nash equilibrium whose cost is unbounded compared to the social optimum. Even for  $m = 2$  machines, we show the same holds true so long as at most  $n/2 - 1$  jobs follow the advertised behavior.
- For party affiliation games [Christodoulou et al. 2006] where players have degree  $\omega(\log n)$  we show a threshold property: any value of  $\alpha < 1/2$  is not sufficient to improve the equilibrium beyond the  $\Omega(n^2)$  price of anarchy, while any value  $\alpha > 1/2$  is sufficient to produce behavior within an  $O(1)$  factor of optimal.

In [Balcan et al. 2010], we consider a more adaptive *learning-based model* where rather than having separate receptive and non-receptive players, each player  $i$  individually decides on each round between following the advertised strategy or acting in a best-response manner with some probability  $p_i$ . Based on their observations, players may then slowly adjust their probabilities over time using some learning rule of their choosing. The goal again is to show that (in expectation or with high probability) this process results in a low-cost state, with as mild as possible conditions on the learning rules used by the individual players. What we show is that requiring only smooth learning updates and a reasonably balanced start (each of the two abstract actions—best response or following the advertisement—initially having a non-negligible weight for each player), we can get a number of positive results. In particular, for fair cost sharing we show that as long as the number of players of each type is sufficiently high, there exists a (polynomially-bounded) time  $T_0$  such that with high probability, for all  $T \geq T_0$ , the cost at time  $T$  is within a logarithmic factor of optimal. In addition, for consensus games, we show that all the positive results for the basic advertisement model carry over to this setting as well, and even better that the minimum degree requirement can actually now be removed. We also analyze a batch version of this learning model where players operate in two distinct phases: an exploration phase (players try each of the two options with equal probability) and an exploitation phase (after some polynomial number of steps the players commit to a specific action that they play from then on). For this we show additional positive results including more general party-affiliation games as well as more general results for fair cost-sharing.

### 3. DISCUSSION AND CONCLUSIONS

Overall, our results show that for a number of important games that have a wide gap between the qualities of their best and worst equilibria, it is possible for a central authority to “nudge” behavior into a high-quality equilibrium via an advertising strategy that only a fraction of players respond to. Moreover, common games

exhibit an interesting structure with some having a smooth tradeoff, some having a threshold property, and some not being amenable to such advertising at all. More broadly, this analysis, and especially our learning-based model, can be viewed as a step to understanding how self-interested learning by the individual players in a game, together with some knowledge of the game's overall structure, can be used to efficiently reach better equilibria.

*Acknowledgements.* This research was supported in part by NSF grant CCF-0953192, ONR grant N00014-09-1-0751, and AFOSR grant FA9550-09-1-0538.

#### REFERENCES

- ANSHELEVICH, E., DASGUPTA, A., KLEINBERG, J., TARDOS, E., WEXLER, T., AND ROUGHGARDEN., T. 2004. The Price of Stability for Network Design with Fair Cost Allocation. In *45th Annual IEEE Symposium on Foundations of Computer Science*.
- BALCAN, M.-F., BLUM, A., AND MANSOUR, Y. 2009. Improved Equilibria via Public Service Advertising. In *ACM-SIAM Symposium on Discrete Algorithms*.
- BALCAN, M.-F., BLUM, A., AND MANSOUR, Y. 2010. Circumventing the Price of Anarchy: Leading Dynamics to Good Behavior. In *Symposium on Innovations in Computer Science*.
- CHRISTODOULOU, G., MIRROKNI, V. S., AND SIDIROPOULOS., A. 2006. Convergence and approximation in potential games. In *23rd Symposium on Theoretical Aspects of Computer Science*.
- CORREA, J. R., SCHULZ, A. S., AND STIER-MOSES, N. E. 2004. Selfish routing in capacitated networks. *Mathematics of Operations Research* 29(4), 236 – 259.
- EVEN-DAR, E., KESSELMAN, A., AND MANSOUR, Y. 2007. Convergence time to Nash equilibrium in load balancing. *ACM Transactions on Algorithms* 3, 3.
- KOUTSOPIAS, E. AND PAPADIMITRIOU, C. H. 1999. Worst-case equilibria. In *Proceedings of the 16th Annual Conference on Theoretical Aspects of Computer Science*.
- NISAN, N., ROUGHGARDEN, T., TARDOS, E., AND VAZIRANI, V., Eds. 2007. *Algorithmic Game Theory*. Cambridge.