

Extended abstract:

Tacit collusion by price algorithms under a
multinomial logit model with demand learning

Thomas Loots and Arnoud V. den Boer
University of Amsterdam

Virtual Conference on Algorithmic Collusion
December 14th, 2020

Motivation. Today, the selling prices of many products and services are determined by algorithms designed to learn which selling prices are in the best interest of its user. There are concerns that algorithms of different firms could learn to *collaborate* instead of compete with each other, in particular in the so-called hub-and-spoke scenario, where different firms implement the same type of price algorithm. Such collaboration or collusion can be detrimental for consumers as it may lead to higher selling prices. However, if the algorithms do not engage in illicit forms of mutual communication to establish the collaboration, legal experts agree that this form of collusion would most likely not be in violation of existing antitrust law.

Whether algorithms are capable of doing this in realistic market scenarios is an open question, and is contested among economists and legal scholars. There is evidently a great urgency to answer this question, as legal experts argue that existing legislation may have to be modified due to the threat of algorithmic collusion. The goal of this paper is to contribute to this debate by showing in a theoretical setting that tacit algorithmic collusion is possible.

Contributions. We study the dynamics of a price-setting duopoly with incomplete information on how demand depends on price. We adopt the widely-used multinomial-logit framework to model demand, and show that straightforward joint-profit maximization may not always lead to sustainable collusion as it does not necessarily increase profits for both firms. We therefore propose an alternative notion of collusion, called fair Pareto-optimal pricing, that ensures equal relative gains compared to profits under the Nash equilibrium for both firms. We show that there exists a unique fair Pareto-optimal price pair that leads to both higher profits and higher selling prices of firms, so that this form of collusion is always profitable for both firms but detrimental for consumers.

Next, we propose a price algorithm that, if playing against itself, learns the fair Pareto-optimal price from accumulating sales data. The algorithm is based on estimating the unknown parameters of the demand model using maximum-likelihood estimation, computing the corresponding estimated fair Pareto-optimal price, and applying an appropriate price perturbation to ensure that prices are sufficiently dispersed to guarantee consistency of the parameter estimates. We show that our algorithm, if deployed by both firms in the duopoly, learns to collude, and derive bounds on the rate of convergence.

Our algorithm can operate in a setting where prices are public but observed demand is private information. Unobservable demand of competing firms poses in principle a non-trivial challenge to the goal of achieving sustainable collusion: not exactly knowing the data set of your competitor hampers statistical estimation of the demand function, and also makes it difficult to check whether deviations by the competitor from the (estimated) collusive price are caused by statistical errors, price experiments, or by a willingness to ‘cheat’ and stop colluding. We show that this challenge can be overcome by noting that the prices charged by the firms *convey information* about their most recent demand observation. By continuously reverse-engineering this information one can completely reconstruct the demand observations of the competitor.