

AI-Driven Prices in Markets

Supervisor: Panayiotis Danassis

November 8, 2021

1 Project Overview & Goal Description

This project aims to develop a practical approach to computing market prices and allocations via a *deep reinforcement learning policymaker* agent, that would allow us to *tune* the prices with regard to diverse objectives such as sustainability and resource wastefulness, fairness, buyers' and sellers' welfare, etc.

In a competitive market, there are numerous producers (sellers) that 'compete' with one another in providing goods and services for consumers (buyers). The theory of competitive markets is one of the most prominent economic models of the 20th century. A fundamental concept in competitive markets is the *market equilibrium* – a stable outcome in which supply equals demand, and all participants are maximally satisfied by the bundles of goods that they buy or sell. The market equilibrium is meant to capture the outcome of a free market, dictated by the market's 'invisible hand', which adjusts the prices until market clearance is achieved via a continuous adjustment of supply and demand (*tâtonnement* process).

Classic market equilibrium results assume that agents can not influence the prices of the markets (i.e., they are price-takers). A great example of a competitive market is farming. There are thousands of farmers and a single farmer can not influence the market or the price based on how much they grow. All the farmer can do is grow the crop and accept the current price for that crop.

In reality, there have been several examples, especially in markets with a limited number of sellers, where instances of collusion and price manipulation have been observed;¹ practices which are known to deter the market from its intended equilibrium outcome. Furthermore, it has been shown that in fundamental market models, the convergence of the *tâtonnement* process is highly dependent on several initial parameters and can therefore be slow. Most importantly, the centralized computation of market equilibria for certain markets can be computationally hard, thus in many cases *impractical* to compute.

Additionally, the market equilibrium is geared towards very specific goals, namely fairness for the participants and economic efficiency given the set of chosen prices. However, from a more global perspective, there are several other important objectives: for example, one of the most pressing issues in modern societies is *sustainability* and the preservation of the Earth's natural resources.² Clearly, the extent to which natural

¹E.g., see <http://news.bbc.co.uk/1/hi/business/7132108.stm> or <https://fortune.com/2015/06/30/apple-conspired-with-book-publishers-appeals-court-confirms/>.

²E.g., see the UN's sustainable development goals <https://www.un.org/sustainabledevelopment/>.

resources are harvested is correlated with the potential monetary gains that the sellers of those resources can achieve in the market. With these ‘exogenous’ objectives being of paramount importance, it is only natural to assume that some form of intervention to the reign of free markets is not only inevitable, but also fully justified.

In this project, we will take an alternative approach, and use deep reinforcement learning to develop a practical approach to computing market prices and allocations, allowing us to optimize a host of diverse objectives such as sustainability and resource wastefulness, fairness, and buyers’ and sellers’ welfare.

2 Project Steps

- Get acquainted with related work (see [1, 2]).
- Implement an approximation algorithm for computing market equilibria in general markets [2] to use as a baseline.
- Implement a deep reinforcement learning agent for computing market prices and allocations, as in [1].
- Perform a detailed empirical evaluation of the policymaker agent.

3 Required Skills

Good programming skills are required (proficiency in Python).

Familiarity with (being able to read and understand) constraint optimization problems (e.g., see [2]).

Being passionate about the topic and good English skills are a must.

References

- [1] P. Danassis, A. Filos-Ratsikas, and B. Faltings, “Achieving diverse objectives with ai-driven prices in deep reinforcement learning multi-agent markets,” *ArXiv: 2106.06060*, 2021. [Online]. Available: <https://arxiv.org/abs/2106.06060>
- [2] R. Garg and S. Kapoor, “Auction algorithms for market equilibrium,” *Mathematics of Operations Research*, vol. 31, no. 4, pp. 714–729, 2006. [Online]. Available: <https://www.cse.iitd.ac.in/~rahulgarg/Publications/2004/GK04.stoc.pdf>