**Internship Title:** Learning User Preferences against Deceptive Users
**Internship Duration:** 6 months

A ***PhD position*** is potentially available after the internship (optional)

**Collaborating Institution**: Massachusetts Institute of Technology (MIT): Prof. Moshe Ben-Akiva; Technical University of Denmark (DTU): Assoc. Prof. Ravi Seshadri; IPP: Andrea Araldo

**Preferred Start Date:** Any **Place**: Palaiseau Campus

**Description of the Internship Project:**

We consider a regulator willing to encourage the sustainable behavior of users, who are confronted with different options when choosing goods, foods, services and mobility modes. To this aim, the regulator can apply appropriate "signals" to users, which may be positive (incentives, subsidies) or negative (prices, taxes, bans). Personalized policies adapt the signal to the needs and preferences of each agent.

To implement a personalized policy, the regulator has to learn the preferences of users by observing their previous choices. Up to now, personalized policies have relied on the hypothesis that users are rational and honest, making each choice in order to maximize their utility. This assumption does not hold in the case of personalized policies, where agents may adopt a deceptive behavior, acting in order to hide their true preferences, with the aim to manipulate the regulator and get a more favorable signal than they would deserve.

The aim of this internship is to analyze the robustness of personalized policies to deceptive agents. We will build upon recent literature in the AI community about strategic interaction.[1-7]

We will model the behavior of users via Random Utility Theory[8] (see Nobel Prize for Economics in 2000), We will model preference elicitation,[9] i.e., the learning process of a regulator running inference methods to learn user preferences. We will model the interaction between agents and the regulator via Stackelberg games.

If users want to deceive the regulator, they need to make choices that do not correspond to their preferred ones, thus paying a "cost of deception". We analytically find under which circumstances such a cost is large enough to deter deceptive behavior. This findings can find regulators to design demand-management policies that are robust to deceptive agents.

**Skills and Qualifications Required:**
Excellent analytic and modeling skills, as well as good programming skills. Previous knowledge of Game Theory and Statistical Methods is a big plus (although not strictly required).

**Main Internship Supervisor:** Assoc. Prof. Andrea Araldo <araldo@telecom-sudparis.eu>

**To apply**: Please send all your notes of your studies at BSc and MSc levels.

## References

[1] Nguyen, T. H.; Wang, Y.; Sinha, A.; and Wellman, M. P. 2019. Deception in finitely repeated security games. In Proceedings of the **AAAI Conference on Artificial Intelligence**, volume 33, 2133–2140

[2] Birmpas, G.; Gan, J.; Hollender, A.; Marmolejo, F.; Rajgopal, N.; and Voudouris, A. 2020. Optimally deceiving a learning leader in stackelberg games. **Advances in Neural Information Processing Systems (NeurIPS)**, 33: 20624–20635

[3] Dawkins, Q.; Han, M.; and Xu, H. 2021. The Limits of Optimal Pricing in the Dark. **Advances in Neural Information Processing Systems (NeurIPS)**, 34: 26649–26660.

[4] Gan, J.; Guo, Q.; Tran-Thanh, L.; An, B.; and Wooldridge, M. 2019a. Manipulating a Learning Defender and Ways to Counteract. In Wallach, H.; Larochelle, H.; Beygelzimer, A.; d'Alché-Buc, F.; Fox, E.; and Garnett, R., eds., **Advances in Neural Information Processing Systems (NeurIPS)**.

[5] Sessa, P. G.; Bogunovic, I.; Kamgarpour, M.; and Krause, A. 2020. Learning to Play Sequential Games versus Unknown Opponents. In Larochelle, H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., **Advances in Neural Information Processing Systems (NeurIPS)**

[6] Nedelec, T.; Calauzenes, C.; Perchet, V.; and Karoui, N. E. 2020. Robust Stackelberg buyers in repeated auctions. In Chiappa, S.; and Calandra, R., eds., Proceedings of the Twenty Third International **Conference on Artificial Intelligence and Statistics**, volume 108 of Proceedings of Machine Learning Research, 1342–1351. **PMLR**

[7] Nguyen, T.; and Xu, H. 2019. Imitative Attacker Deception in Stackelberg Security Games. In **International Joint Conferences on Artificial Intelligence (IJCAI)**, 528–534.

[8] Strzalecki, T. 2025. Random Utility, 3–18. Econometric Society Monographs. Cambridge University Press, link

[9] Chapman, J.; and Fisher, G. 2025. Preference Elicitation: Common Methods and Potential Pitfalls. In Snowberg, E.; and Yariv, L., eds., Handbook of Experimental Methodology, Volume 1. Chapter 2.