

Research Statement

Emmanouil Pountourakis

My research interests lie in the intersection of theoretical computer science, economics, and game theory. The main focus of my work has been in algorithmic mechanism design, the theory of algorithm design in strategic environments. In my earlier research, I studied the performance loss of optimal algorithms in the presence of incentive and computational constraints in the areas of group-strategyproof cost-sharing [11, 6], stable marriage [2], and mechanism design without money [10].

In these explorations I have been intrigued by the fact that despite how well several aspects of mechanism design are understood, theoretically optimal mechanisms are rarely observed in practice. This phenomenon, the misalignment of predicted and observed behavior, also occurs in the participants of the mechanism. Recently I have focused on understanding the observed practices and providing theoretical understanding for their prevalence.

In the following I outline my three main current research directions towards this goal.

Simplicity and approximation

Traditional mechanism design focuses on designing mechanisms (algorithms) that optimize a given objective subject to incentive constraints. However, theoretically optimal mechanisms rarely manifest in practice. Instead, we observe simple mechanisms that are intuitive to participants. One of the main contributions of algorithmic game theory has been the introduction of the theory of approximation to mechanism design: what is the optimality trade-off of mechanisms that satisfy other important properties like simplicity, determinism, or robustness?

Myerson [8] characterized the revenue maximizing auction in a single item setting. A surprising corollary of this result is that the optimal auction is deterministic even if we allow for randomized allocations, and, if we additionally assume that agents are a priori identical, the widely prevalent second price auction with anonymous (common) reserve is optimal. These consequences of Myerson's result are perhaps the reason why it is one of the most celebrated results in microeconomics.

However, most real life applications are inherently asymmetric due to public bidding history and a variety of demographic information available to many auction platforms, e.g. eBay. Yet second price with anonymous reserve and anonymous posted pricing are by far the most prevalent auctions used in practice.

A natural question is to study the approximation factor of these mechanisms. This was partially resolved by [5] showing that the approximation factor of second price auction with anonymous reserve lies in $[2, 4]$. Recently, we provide a tight analysis of the approximation ratio of anonymous pricing to the standard upper bound benchmark of ex-ante relaxation [1]. This implies tightened upper bounds for the approximation of optimal auction from four to $e \approx 2.718$ for both anonymous posted pricing and second price auction with anonymous reserve providing the first improvement to this open question in the last half decade. Despite our efforts, the approximation ratios of anonymous posted pricing and second price auction with anonymous reserve remain unresolved. In the future, I would like to provide a definite answer to these questions. Some of our findings already suggest that a direct comparison between the second price auction with anonymous reserve and the ex-ante relaxation benchmark is a promising approach and the answer is most likely 2, matching the previous known lower bound.

In the future, I would like to develop similar techniques in order to resolve elusive questions in more complicated settings in mechanism design. For instance, revenue maximization in multi-dimensional settings is either poorly understood or inefficient to compute. Consequently, it becomes imperative to study the approximation ratio of simple auctions. One such example is selling multiple items to a unit-demand bidder. Even when the values of the buyer for the different items are distributed independently the optimal auction

may be randomized. Several open problems of approximation arise in this setting, e.g. the approximation ratio of deterministic mechanisms [3]. I am particularly interested in resolving the approximation ratio of uniform pricing to the optimal randomized allocation, for which our result [1] together with Chawla et al [3] imply an upper bound of $2e$.

One of the main challenges in our tight analysis of comparing anonymous pricing to the ex-ante relaxation benchmark is that the worst case approximation ratio is realized in the limiting case where the number of buyers approaches infinity. Generally, algorithmic mechanism design has many similar unresolved questions of constant approximation, where the worst case instance is realized in the limit. I believe that the tools we developed for analyzing the approximation of anonymous posted pricing will be useful in designing a theoretical framework to tackle the open problems I outlined earlier and other similar questions.

Dynamic environments

Most of the contributions of algorithmic game theory has been in the study of static environments, where different entities interact in a single-shot game fashion. However, most realistic scenarios involve several interactions across time. For instance, advertisers bid over and over again for display ads. In such situations, the assumption that strategic agents behave myopically can be far from the truth. For example, the advertisers are likely to take into account their future pay-offs as well when deciding on their bids. Hence, it becomes imperative to study algorithmic design in a dynamic setting where incentives take future interactions into consideration.

One of my current projects in this area studies the problem of repeated sales of a digital good over several periods. Our goal is to study the revenue of the seller in a perfect Bayesian equilibrium, i.e. if the seller has no commitment power. Unfortunately, there are several negative results for the single buyer version of this problem [9, 12, 4] which show that the seller's revenue can be arbitrarily low. In comparison, a seller with commitment power can generate the optimal revenue of the static variant of the problem in each period. The negative implications of the absence of commitment power carry over to the multiple buyer case if the seller is allowed to use an arbitrary pricing scheme, because then the problem reduces to several independent single buyer cases.

Obviously this phenomenon does not occur in practice. One possible explanation may be that the seller cannot choose prices arbitrarily. For example, discriminatory pricing may be considered unethical or, in some circumstances, illegal. Motivated by this observation we address this problem by restricting the seller to use anonymous pricing schemes. Our preliminary results show that in several cases this restriction of the seller's action space leads to increased revenue compared to the unrestricted case. We are currently generalizing these results by studying how this revenue difference evolves as the set of buyers increases. Our goal is to provide approximation guarantees for the revenue compared to a seller with commitment power as a function of the number of buyers.

Our preliminary result demonstrates that dynamic settings can be far more complicated than their static counterparts. There are two important phenomena that are highlighted through this project. First, allowing the mechanism designer to optimize in every single period may harm the overall optimization of the given objective without commitment power. Secondly, limiting cases, e.g. a single buyer, yielded negative results that were unrealistic. Our positive results were obtained after increasing the number of buyers. This suggests that large market assumptions are imperative to developing theories with predictive power in dynamic environments. In my future research, I would to extend these observations in the analysis of more complicated dynamic environments, for example, repeated sales with a limited supply of items and the buyers have to compete.

Behavioral models

One of the major assumptions in game theory is that strategic agents are perfectly rational utility maximizers. However, such behavior is rarely observed in practice. Agent behavior often deviates from theoretical predictions as the complexity of the mechanism or the environment increases.

In dynamic environments this issue is exacerbated due to the added element of decision making over time. The standard model employed in economics is geometric discounting. According to this model agent

behavior is time-consistent, i.e. the agent will remain consistent with their past choices. However, time-inconsistency is clearly inherent to human behavior. A large literature in behavioral economics addresses this issue through the notion of hyperbolic discounting. The simplest version of this is present bias: when making a decision a present biased agent overestimates today’s payoff by a multiplicative factor. In a recent paper, Kleinberg and Oren [7] model the behavior of a present biased agent as traversing a graph in order to reach from a source to a destination. They study the cost of procrastination: the ratio of the cost of the path followed by a present-biased agent to the shortest path that connects the source and the destination. They characterize the cost of procrastination using the underlying properties of the graph.

In a recent paper, we study this problem from a mechanism design perspective: A creditor seeks to design a contract for a debtor, that seeks to pay-off her debt. We propose a Bayesian version of the aforementioned behavioral model for the debtor, where her present bias factor is drawn from a distribution in each period. The creditor has access to this distribution and seeks to exploit the information to maximize revenue. We solve the optimization problem of the creditor by drawing an interesting connection with pricing theory. In addition, we propose several regulations on the contract design to reduce the creditor’s exploitative power.

In my future research, I would like to incorporate observations of Behavioral Economics broadly in mechanism design. There is a large literature that tries to understand the notion of consumer fairness where consumers react adversely to price discrimination. Studying mechanism design under a behavioral model of consumer fairness may provide an additional justification for why non-discriminatory mechanisms, like anonymous posted pricing, are widely prevalent in practice.

Conclusion

Through my research I want to provide a theoretical justification for the simple mechanisms we observe in practice. Towards this goal, I have studied approximation in mechanism design, analyzed dynamic settings, and incorporated behavioral assumptions. In addition, these approaches work in synergy towards my main objective: approximation is necessary to understand mechanisms in dynamic settings and perfectly rational decision making in a dynamic environment is not only unrealistic it is also what drives several of the negative results.

References

- [1] Saeed Alaei, Jason D. Hartline, Rad Niazadeh, Emmanouil Pountourakis, and Yang Yuan. Optimal auctions vs. anonymous pricing. In *Proceedings of the Fifty-Sixth Annual IEEE Symposium on Foundations of Computer Science*, FOCS ’15, 2015.
- [2] Georgios Askalidis, Nicole Immorlica, Augustine Kwanashie, David Manlove, and Emmanouil Pountourakis. Socially stable matchings in the hospitals/residents problem. In *Algorithms and Data Structures - 13th International Symposium, WADS 2013, London, ON, Canada, August 12-14, 2013. Proceedings*, pages 85–96, 2013.
- [3] Shuchi Chawla, David L. Malec, and Balasubramanian Sivan. The power of randomness in bayesian optimal mechanism design. *Games and Economic Behavior*, 91:297–317, 2015.
- [4] Nikhil R. Devanur, Yuval Peres, and Balasubramanian Sivan. Perfect bayesian equilibria in repeated sales. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA ’15, pages 983–1002. SIAM, 2015.
- [5] Jason D. Hartline and Tim Roughgarden. Simple versus optimal mechanisms. *SIGecom Exchanges*, 8(1), 2009.
- [6] Nicole Immorlica and Emmanouil Pountourakis. On budget-balanced group-strategyproof cost-sharing mechanisms. In *Internet and Network Economics - 8th International Workshop, WINE 2012, Liverpool, UK, December 10-12, 2012. Proceedings*, pages 244–255, 2012.

- [7] Jon M. Kleinberg and Sigal Oren. Time-inconsistent planning: a computational problem in behavioral economics. In *ACM Conference on Economics and Computation, EC '14, Stanford , CA, USA, June 8-12, 2014*, pages 547–564, 2014.
- [8] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58–73, 1981.
- [9] Jean Tirole Oliver D. Hart. Contract renegotiation and coasian dynamics. *The Review of Economic Studies*, 55(4):509–540, 1988.
- [10] Emmanouil Pountourakis and Guido Schäfer. Mechanisms for hiring a matroid base without money. In *Algorithmic Game Theory - 7th International Symposium, SAGT 2014, Haifa, Israel, September 30 - October 2, 2014. Proceedings*, pages 255–266, 2014.
- [11] Emmanouil Pountourakis and Angelina Vidali. A complete characterization of group-strategyproof mechanisms of cost-sharing. *Algorithmica*, 63(4):831–860, 2012.
- [12] Klaus M. Schmidt. Commitment through incomplete information in a simple repeated bargaining game. *Journal of Economic Theory*, 60(1):114 – 139, 1993.