Continuous Mechanism Design

Ruben Juarez
Department of Economics, University of Hawaii
rubenj@hawaii.edu

August 30, 2018

Abstract

Despite the accumulating literature on mechanism design over the last decades, the equilibrium outcome of most mechanisms described to date fails to be robust to small changes in the specifications and parameters of the model. Lacking such robustness, the designer may fail to implement his desired outcome. Thus, here we argue that new mechanisms should be more resistant to physical, modeling and behavioral misspecifications.

Over the last decades, the field of mechanism design has shown that incentives can achieve desirable objectives when players are rational. This has been validated in a variety of applications ranging from auctions and cost-sharing to matching and voting. The natural engineering approach of the field to construct such mechanisms requires the designer to make strong assumptions about the preferences and the behavior of the agents, in particular. However, errors can and do occur. The designer might be incorrect about the parameters applied to the model, the preferences of the agents, or their behaviors. Even if the designer was initially correct, agents may change their preferences or alter their behavior over time. Traditional mechanism design is not immune to these changes, including those that may be subtle. Therefore, there is a need in the field to design robust mechanisms that would, in effect, be error-proof.

While it might not be realistic to design completely resilient mechanisms that work for every possible change, our goal is to study the subtle changes that might happen due to misestimation of the parameters, preferences, or behavior of the agents to minimize potential errors. We posit that such *small changes should not have a large impact on the outcome (equilibrium) of the mechanism*. To formalize this, continuity requires that sufficiently small changes in the input of a mechanism (parameters, preferences, etc.) result in arbitrarily small changes in the output (payoff of the agents, the objective of the planner, etc.). Continuity in the outcome of the mechanism is a desirable property, as the parameters, preferences, and behavior of the agents are hard to precisely estimate, and small misestimations should not have a significant impact on the outcome of the mechanism. The lack of continuity in the mechanism has implications for efficiency and fairness, both of which are central objectives in mechanism design.

It is important to clarify that we focus here on outcome continuous mechanisms, which take into account the final allocation to the agents at equilibrium. Because the (presumably Nash) behavior of the agents determines the allocation of the agents, this continuity property is substantially stronger than traditional continuity requirements and has received little attention in the literature.[1]

The study of continuous mechanism design fits within the extensive literature of robust (prior-free) mechanism design, where the intent is to design mechanisms that are robust for an uninformed planner about the characteristics of agents participating in the mechanism — e.g., when agents know each other and may collude ([20, 23, 12]) or when the planner does

[1] An exception is the literature on payoff-continuity when there is incomplete information in a game, see [17, 24, 18]. Unfortunately, such a literature have had limited applicability to specific mechanism design problems.

not have a priori knowledge about the utility of the agents ([21, 11, 19])[2]. Unlike most of the work previously done in robust mechanism design, we propose to examine and study the impact of small changes to the resiliency of the model.[3]

The complexity of engineering resilient mechanisms creates difficult challenges for the designer. To illustrate this, we focus on specific mechanism design settings applied to cost-sharing problems, however, similar examples can be extended to more general mechanism design settings where money is available.

Continuity with respect to the parameters in the model

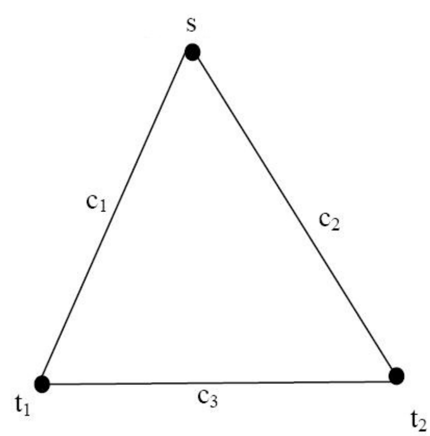

Figure 1: Network with two agents at the common source s and two different sinks t_1 and t_2.

Consider two agents, 1 and 2, interested in connecting from the source s to t_1 and t_2, respectively (see figure 1). The agents simultaneously choose paths in the network to connect their unique source to their unique sink, and the choice of paths of all the agents leads to the network that must be constructed and/or maintained. A mechanism splits the total cost of the network formed between the two agents who care about minimizing their cost-share.

The Shapley mechanism distributes the cost of the link selected equally among the agents. This mechanism is considered fair, is easy to implement (decentralized), and has been

[2]Due to space limitations, we omit most of the extensive literature about the worst-case analysis from the fields of Economics and Computer Sciences.

[3]A natural open question is to understand to what degree local changes lead to global properties. A recent strand of literature has focused on whether local strategy-proofness implies global strategy-proofness ([3, 25, 2])

widely studied in the literature ([1, 4, 5, 6, 7]). Unfortunately, the equilibrium of the Shapley mechanism is not continuous when simple measurement errors in the cost of the links occur. Indeed, consider the case of a symmetric problem where the costs are $c_1 = c_2 = c > 0$ and $c_3 = \frac{c}{2} + \epsilon$, where ϵ is a small error.

When $\epsilon > 0$, the unique Nash equilibrium requires each agent to connect directly, select st_1 and st_2, and pay a cost equal to c. However, when the error term ϵ is negative, $\epsilon < 0$, there are exactly two Nash equilibriums (st_1, st_1t_2) and (st_2t_1, st_2), both of which require the agents to coordinate. The payoffs to the agents are $(\frac{c}{2}, c + \epsilon)$ or $(c + \epsilon, \frac{c}{2})$, respectively.

This discontinuity of the Nash equilibrium is highly undesirable in many dimensions. First, the payments of the agents substantially change, creating highly asymmetric payoffs when the entire problem is symmetric. Second, the strategies of the agents widely differ depending on whether ϵ is positive or negative. For a positive ϵ, the agents' strategies are simple since they do not require coordination. In such a case, agents pay the cost of connecting directly to their link. However, when ϵ is negative, agents need to coordinate on the route selected, thus creating added complexity in the strategies.

While there are some recent characterizations regarding the continuity of the Nash equilibrium in a variety of cost-sharing models, mainly providing highly restrictive mechanisms that only depend on the aggregate cost and disregard the demand of the agents[4], more work is needed to address the role of payoff-continuity in more general cost-sharing and allocation problems. The study of continuity beyond payoffs, for instance by focusing on "similar" strategies, warrants further examination.

Continuity with respect to the preferences of the agents

Most economic theory assumes that individual behavior is driven by rational decisions, with a few theories challenging this approach. We believe that many failures of the theory of mechanism design occur not due to a rationality assumption, but due to a design failure, such as eliciting appropriate information from the agents. Indeed, we find two severely under-explored issues that can better link the theory with the practice of mechanism design.

Preference Representation: Humans are complex and so are their preferences. Representing preferences of individuals is particularly difficult in cost-sharing problems, where agents

[4]See, [14] for payoff-continuity in a routing model, [13] for payoff-continuity in a sequential sharing model, and [16] for continuity in a time allocation model. [9, 10] also characterize cost minimizing networks when the announcement rule can change.

might have preferences over a continuum of options (e.g., the division of an infinitely divisible resource among agents). It is the designer's responsibility to elicit information from the agents that approximate their complex preferences, for instance by eliciting simple information (e.g., a one parameter function [22, 8]) or a more accurately description of their taste.[5]

A mechanism that solicits complex preferences from the agents will likely produce large changes in the final allocation of the agents at equilibrium (especially when agents have difficulty estimating their own preference), which in turn might create even more complexities to predict the equilibrium. Therefore, a trade-off for the designer arises: soliciting simple information from the agents that is more stable to changes yet potentially farther away from the planner's objectives (e.g., efficiency, fairness, etc.) versus soliciting more complex information with less predictive power yet closer to the planner's objectives when the prediction actually works.

Externalities: Perhaps one of the largest under-explored issues in mechanism design occurs when agents experience externalities. Such externalities arise when agents care not only about their own utility (i.e., allocation, cost-share, etc.) but also about that of others. Indeed, in the presence of social networks, agents may experience a satisfaction (or dissatisfaction) when their close friends (or enemies) are better off. A particular case of externalities occurs when agents experience envy, i.e., the preference of someone else's allocation to his own. However, in general, the externalities experienced by agents are hard to estimate and model into the utility function.

The study of externalities presents a challenge for the designer, who might be uninformed about the social network and its magnitude in the utility of the agents. Therefore, there is a need to understand how new mechanisms will change in the presence of externalities, and identify more resilient mechanisms that work well under such circumstances.[6]

[5] There is also a large literature on what preferences are in the social sciences, and more importantly, the oversimplification of preferences and behavior in Economics. Both oversimplification limit the applicability of Economics and cross-talk among fields.

[6] This happens even in the most simplistic model of resource sharing of a fixed unit of a good, where [15] shows the difficulty to tackling this problem when the planner is uninformed about the social network and its effect on the utility function. It characterizes a class of strategy-proof resource allocating mechanisms.

Is the Nash equilibrium the right equilibrium?

The predictive power of the Nash equilibrium has been widely tested, especially in the experimental literature, results of which have been mixed. The success of prediction is specific to every problem and highly dependent on the type of equilibrium that it generates; e.g., there are positive results for games that generate a strong Nash equilibrium, and even some games with Pareto dominant equilibria. The accuracy of the Nash equilibrium is also dependent on several other exogenous factors, such as the framing of the problem or the type of agents who are playing.

As theorists, our research should be informed by the experimental literature as an aid to refine our models to more closely resemble what has been observed in real-world settings. This is particularly crucial in the mechanism design literature, where the planner often has a menu of mechanisms, each generating equilibrium(s) with different properties, some of which may or may not appear in reality.

Specifically, consider the example represented by figure 1 when $\epsilon > 0$. The Shapley mechanism produces a unique inefficient equilibrium where each agent pays c, for a total cost of $2c$, whereas the cheapest way to connect the agents costs $\frac{3c}{2} + \epsilon$. A simpler cost-allocating mechanism that divides the total cost of the network formed in proportion to the cheapest connection of their source-sink (in this case c for each agent) generates two equivalent equilibrium, where each agent pays $\frac{3c}{4} + \frac{\epsilon}{2}$. Similarly to Shapley, this mechanism might produce multiple equilibria, but all of them will be Pareto ranked. Thus, if the Nash equilibrium is a good predictor of the outcome implemented by the mechanism, the least inefficient equilibrium stands out as a reasonable selection that all the agents prefer ([14]). This produces a higher predictive power than Shapley, even when the costs c_1, c_2 and c_3 change slightly.

Conclusion

Mechanism design has been a major component in the economic analysis of institutions and markets, especially when the designer has incomplete information about the participants. While there has been a recent expansion in the field with interest in robust mechanisms that work well when the planner has little information about the agents, the field can be improved by identifying the right trade-offs that account for physical, modeling and behavioral errors, some of which might be alleviated at the design stage and by better linking the theory with more practical studies.

References

[1] E. Anshelevich, A. Dasgupta, L. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. *In 45th Annual IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 59–73, 2004.

[2] A. Archer and R. Kleinberg. Truthful germs are contagious: a local-to-global characterization of truthfulness. *Games and Economic Behavior*, 86:340–366, 2014.

[3] G. Carroll. When are local incentive constraints sufficient? *Econometrica*, 80(2):661–686, 2012.

[4] H. L. Chen, T. Roughgarden, and G. Valiant. Designing networks with good equilibria. *Proceedings of the 19th annual ACM-SIAM Symposium on Discrete Algorithms*, 2008.

[5] A. Epstein, M. Feldman, and Y. Mansour. Strong equilibrium in cost sharing connection games. *ACM conference on Electronic Commerce, San Diego*, 2007.

[6] A. Epstein, M. Feldman, and Y. Mansour. Equilibrium in cost sharing connection games. *Game. Econ. Behav.*, 67:51–68, 2009.

[7] A. Fiat, H. Kaplan, M. Levy, S. Olonetsky, and R. Shabo. On the price of stability for designing undirected networks with fair cost allocations. *ICALP06*, 2006.

[8] A. V. Goldberg and J. D. Hartline. Collusion-resistant mechanisms for single-parameter agents. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 620–629. Society for Industrial and Applied Mathematics, 2005.

[9] J. L. Hougaard and M. Tvede.
Truth-telling and nash equilibria in minimum cost spanning tree models.
European Journal of Operational Research, 222(3):566–570, 2012.

[10] J. L. Hougaard and M. Tvede.
Minimum cost connection networks: Truth-telling and implementation.
Journal of Economic Theory, 157:76–99, 2015.

[11] R. Juarez.
The worst absolute surplus loss in the problem of commons: random priority versus average cost.
Economic Theory, 34(1):69–84, 2008.

[12] R. Juarez.
Group strategyproof cost sharing: the role of indifferences.
Games and Economic Behavior, 82:218–239, 2013.

[13] R. Juarez, C. Y. Ko, and X. Jingyi.
Sharing sequential values in a network.
Journal of Economic Theory, 177:734–779, 2018.

[14] R. Juarez and R. Kumar.
Implementing efficient graphs in connection networks.
Economic Theory, 54(2):359–403, 2013.

[15] R. Juarez and R. Kumar.
Dividing a dollar in a network.
2018.

[16] R. Juarez, K. Nitta, and M. Vargas.
Profit-sharing and efficient time allocation.
Technical report, Mimeo, University of Hawaii, 2018.

[17] A. Kajii and S. Morris.
The robustness of equilibria to incomplete information.
Econometrica: Journal of the Econometric Society, pages 1283–1309, 1997.

[18] A. Kajii and S. Morris.
Payoff continuity in incomplete information games.

journal of economic theory, 82(1):267–276, 1998.

[19] A. Mehta, T. Roughgarden, and M. Sundararajan.
Beyond moulin mechanisms.
In *Proceedings of the 8th ACM Conference on Electronic Commerce*, pages 1–10. ACM, 2007.

[20] H. Moulin.
Incremental cost sharing: Characterization by coalition strategy-proofness.
Social Choice and Welfare, 16(2):279–320, 1999.

[21] H. Moulin.
The price of anarchy of serial, average and incremental cost sharing.
Economic Theory, 36(3):379–405, 2008.

[22] H. Moulin.
One-dimensional mechanism design.
Theoretical Economics, 12(2):587–619, 2017.

[23] H. Moulin and S. Shenker.
Strategyproof sharing of submodular costs: budget balance versus efficiency.
Economic Theory, 18(3):511–533, 2001.

[24] C. G. Rothschild.
Payoff continuity in incomplete information games: A comment.
Journal of Economic Theory, 120(2):270–274, 2005.

[25] S. Sato.
A sufficient condition for the equivalence of strategy-proofness and nonmanipulability by preferences adjacent to the sincere one.
Journal of Economic Theory, 148(1):259–278, 2013.