

A Report on the Workshop on Mechanism Design for Social Good

REDIET ABEBE

Cornell University

and

KIRA GOLDNER

University of Washington

We present a report on the first Workshop on Mechanism Design for Social Good (MD4SG '17), which took place at the 18th ACM Conference on Economics and Computation (EC '17), and discuss unifying themes for research at this interface.

1. INTRODUCTION

The first Workshop on Mechanism Design for Social Good (MD4SG '17) was held in conjunction with the 18th ACM Conference on Economics and Computation (EC '17) at MIT in Cambridge, MA on June 26, 2017. The workshop featured keynote talks, contributed presentations, a poster session, and a panel discussion with researchers from the economics and computation community and adjacent fields.

The EC community has made great strides both in the development of theoretical foundations and in applications of algorithms, optimization, and mechanism design. Key application domains have so far included ad auctions and electronic commerce, fair division, prediction markets, and social networks. The first Workshop on Mechanism Design for Social Good focused on a set of promising applications aimed at *improving access to opportunity*, specifically within housing, healthcare, and education. Examples include low-income housing allocation, refugee resettlement, diversity and integration in education, and aligning incentives in healthcare.

In these domains of interest, the government and individuals design allocation policies, impose regulations on activities, and determine protocols of interaction. Hence, algorithmic and mechanism design insights have the potential to significantly improve societal welfare, especially for communities of people where opportunities are traditionally limited. This requires a deep understanding of the domains of interest in order to formulate problems where more efficient and fair solutions can be implemented to improve access to opportunity. Consequently, the workshop included three keynote presentations by domain experts from the fields of economics and public policy, which we discuss in detail in Section 2. Each talk was followed by a discussion period to foster future learning and collaborations, as well as jointly brainstorm research directions.

The workshop also consisted of ten contributed talks and a poster session highlighting exemplary work falling under this theme. Participants spanned areas including computer science, economics, industrial engineering, operations research,

management sciences, and public policy. We highlight their talks in Section 3.

Working across disciplinary lines, implementing theoretical insights in practice, and developing a framework for applications to inform theoretical work comes with a unique set of opportunities and challenges. To discuss these, share best practices, and formulate future directions for the field, the workshop concluded with a panel discussion with researchers who have extensive experience working at this interface. This is summarized in Section 4.

This workshop grew out of a multi-institutional, interdisciplinary research group with the same objective that the authors co-founded and have been co-organizing since September 2017 [Abebe and Goldner, 2016].¹ Both the workshop and the research group are part of a broader vision to foster interest in research at this interface and increase collaborations across disciplinary boundaries, as well as with practitioners, in order to improve societal welfare.

2. KEYNOTE TALKS

The workshop consisted of three keynote presentations focused on settlement, health-care, and education, with talks by Scott Duke Kominers (Harvard Business School and Harvard Department of Economics), Mark Shepard (Harvard Kennedy School), and Parak Pathak (Massachusetts Institute of Technology, Department of Economics). We present a summary of these talks below:

Scott Duke Kominers: Good Markets (Really Do) Make Good Neighbors. This presentation consisted of a discussion on how to apply market design to improve societal welfare; two case studies focus on refugee settlement and health-care data exchange.

Market design aims to apply economic principles to design and re-design market institutions. This emerges in a variety of domains and presents an opportunity for researchers to contribute to the theory, application, and evaluation of market-design-inspired solutions. This can, in turn, inform research in the theory and applications of market design. Two key dimensions are when the designer can set rules governing what types of transactions may occur in a market and when the designer can build infrastructure for facilitating transactions. Some potential types of interventions inspired by market design are: (1) *marketplace mechanism re-design*, in cases where a market exists, but is not performing satisfactorily, (2) *information provision*, where participants have unequal access to or incentive for collecting information, (3) *re-shaping the extensive margin*, where a marketplace exists but agents either are not participating or wholly lack access to the marketplace, and (4) *market creation*, in cases where a market is missing.

One such opportunity arises in *refugee resettlement*. Each year hundreds and thousands of refugees have to be resettled. Initial resettlement areas matter both for the refugees and for the communities they join. Providing better matches that take into account local communities, labor markets, education, and the skills and preferences of the refugees can improve long-term socioeconomic outcomes, and lead to more supply of resettlement areas. However, current matching processes are rela-

¹To learn more or join the MD4SG email list, please email the organizers at organizers@md4sg.com.

tively ad hoc. Delacrétaz et al. [2016] consider the problem of designing a matching system that incorporates the refugees' and localities' preferences while respecting local service and resource constraints. They develop a model with multidimensional combinatorial constraints and offer mechanisms for different institutional contexts. This work is coupled with Refugees' Say, an organization co-founded by Will Jones and Alex Teytelboym, which aims to implement some of these insights in practice [Jones and Teytelboym, 2016].

A second example is *healthcare data exchanges*. Patients may receive care at multiple institutions, which can lead to duplication of patient health data across sites and errors in reporting health statistics such as the number of deaths and extent of treatment. Sharing health-related data for research is inhibited by the Health Insurance Portability and Accountability Act (HIPPA), which restricts data-exchange in order to protect patients. This challenge presents an opportunity to design a data exchange market which preserves HIPPA protections, while incentivizing participation. Kho et al. [2015] designed and implemented a tool that creates a secure, privacy-preserving linkage of electronic health data across multiple sites in Chicago. This software has been successfully used to link and de-duplicate 7 million records across six institutions and reduce duplication of patient records by as much as 28%.

Mark Shepard: Making Health Insurance Work through Market Design. Mark Shepard's work focuses on health economics, and in particular on health insurance in the United States. In his talk, Shepard summarized the history behind the recent rise of centralized insurance markets in the U.S. In many settings today (e.g. Medicaid, Medicare, and the ACA exchanges), the government creates a market where (subsidized) private health insurance plans are offered to patients. However, the government is currently not fully harnessing its role as market maker. Adverse selection, moral hazard, human failure, and other incentive issues plague insurance markets. Shepard's talk outlined suggestions for how tools from mechanism design could be helpful in re-aligning insurance provider incentives and increasing patient welfare.

One instance is the use of risk-adjustment to mitigate *adverse selection*. In order to make patients with different levels of health risk approximately equally attractive as insurance customers, the government can offer a subsidy to the insurance companies for each patient, which is determined in relation to the patient's expected health costs for the year.

Another example is related to *competitive procurement* of health plans. When an institution, such as a university, is determining what insurance plans to offer their employees, they review several plans from multiple insurance companies and select a few to offer to their employees, presumably choosing the quantity and type of plans in order to best trade off patient health and employer cost. The government can choose to use its role as market maker in a similar way, creating competition to enter the market in order to encourage that prices or quality of plans be set in a more desirable manner.

In summary, U.S. health insurance markets are afflicted with several inefficiencies, which indicate both inadequate alignment of incentives and market failures. Insights from mechanism design can be used to address some of these challenges to create more effective systems that work for individuals, institutions, and insurance

companies alike.

Parag Pathak: Mechanism Design for Education. A canonical example of a research area where mechanism design has been used to implement new solutions aimed at improving societal welfare is school choice. Driven by mechanism design research, over the past couple of decades, there have been significant changes in how students are assigned to public schools across the country. There are a variety of other research directions in education where solutions guided by mechanism design can improve efficiency and increase fairness in the allocation of resources.

One such example is *affirmative action*. Affirmative action, especially in admission to selective educational institutions, is a contentious topic with numerous Supreme Court cases such as the 2003 Seattle/Louisville vs. PICS. Many cases have either impeded the implementation of race-based plans and/or continue to challenge it in many of America's major cities [*Grutter vs Bollinger*, 2003; *Seattle/Louisville vs. PICS*, 2007; *Fischer vs. University of Texas*, 2016]. Race-neutral plans, which provide an alternative to race-based quotas, are being used more widely as a result. A key question is how can we achieve diversity without using a race-based plan? One solution used by Chicago Public Schools (CPS) is a place-based affirmative action system.

Motivated by CPS, Ellison and Pathak [2016] study affirmative action systems with reserves. CPS serves a population of about 400,000 students, about 40% of whom are African American and 45% of whom are Hispanic. It has 10 selective-entry high schools, which get about 16 thousand applications every year. CPS used a race-based allocation scheme before switching to a race-neutral, place-based system in 2008. Ellison and Pathak [2016] explore the efficacy of race-neutral alternatives in substituting racial quotas, as well as discuss trade-offs between allocation efficiency and diversity. They show that if student outcomes depend on how closely curricula are tailored to their ability, then exam-based allocations are desirable. However, if there is a diversity objective, then the optimal allocation system is one that admits applicants based on exam scores and then provides affirmative action to groups with lower scores. They explore different targeting schemes and find that affirmative action systems allow for an applicant to be assigned via multiple routes.

3. CONTRIBUTED TALKS

The workshop also highlighted research by researchers across disciplines through ten contributed talks and a poster session. We present a summary of the contributed talks by Peng Shi (University of Southern California, Marshall School of Business), Benjamin Roth (Harvard Business School), Can Zhang (Georgia Institute of Technology, H. Milton Stewart School of Industrial and Systems Engineering), Nika Haghtalab (Carnegie Mellon University, Department of Computer Science), Hamsa Bastani (IBM Research), Imanol Arrieta-Ibarra (Stanford University, Management Science and Engineering), Yunan Li (University of Pennsylvania, Department of Economics), Vince Conitzer (Duke University, Department of Computer Science), and Hongyao Ma (Harvard University, Department of Computer Science). The workshop website contains references to poster presentations.

Peng Shi: How (Not) to Allocate Affordable Housing. A key issue in addressing inequality is expanding the availability of affordable housing for families in

need. Motivated by the allocation of affordable housing in New York City, Arnosti and Shi [2017] analyze a dynamic assignment problem in which the social planner would like to balance two goals: providing assistance to the most disadvantaged families and providing high-quality matches. They show that these two goals are at odds; in particular, NYC’s current system of repeated lotteries with unlimited entry effectively targets disadvantaged applicants, but achieves poor match quality compared to first-come first-served waiting lists. They also consider alternate mechanisms—allowing agents to use unused tickets, using waiting lists, and limited entry—and show surprising equivalence results between these mechanisms. Finally, they discuss when it is beneficial to prioritize the quality of matches to targeting families in need, and its implications to allocation of affordable housing in practice.

Benjamin Roth: Targeting High Ability Entrepreneurs Using Community Information: Mechanism Design in the Field. Communities hold valuable information about their members, which can be used to improve screening and allocation efforts. Hussam et al. [2017] study one such example: predicting entrepreneurship to provide access to credit. They conduct a large-scale field experiment and demonstrate that community members are able to predict successful micro-entrepreneurs in Maharashtra, India. This information remains useful even after controlling for observables. Since incentive issues arise due to interpersonal relationships, soliciting accurate information is not straightforward. For instance, community members distort their reports in favor of family members and close friends when real grants are being distributed. However, simple techniques such as peer prediction, public reporting, and cross-reporting can be used to effectively realign incentives and elicit highly predictive information about the marginal return of capital.

Can Zhang: Truthful Mechanisms for Medical Surplus Product Allocation. Access to quality medical products is a concern in many health facilities in developing nations. Medical Surplus Recovery Organizations (MSROs) assign surplus medical products found in facilities in developed nations to under-served health facilities in the developing world. However, this assignment process can be inefficient due to the uncertain and uncontrollable nature of both the supply and demand. Zhang et al. [2017] provide a solution inspired by mechanism design to allocate products truthfully and efficiently in a setting where the health facilities are strategic and non-profit. MSROs solicit each recipient’s preference ranking of different products. But, in the generic setting, the only truthful mechanism is one that randomly selects a recipient at each shipping opportunity, and thus does not use this information. It can also lead to very inefficient allocation outcomes. Zhang et al. [2017] find that two strategies—(1) not sharing the inventory, and (2) not sharing information about other recipients—can greatly improve MSROs’ value provision.

Nika Haghtalab: Opting Into Optimal Matching. A seminal mechanism design problem with social good applications is kidney exchange, where an incompatible pair of a kidney donor and a transplant patient is matched with another incompatible donor-patient pair such that when they switch donors, the transplants are both compatible. The impact of kidney exchanges is growing in recent years as hospitals enroll patients into regional or even national kidney exchange programs.

Hospitals, however, may not find it individually rational to participate in the exchanges if they are able to match more of their own patients outside of the exchange, which would reduce the realized social welfare of a mechanism. Blum et al. [2017] consider the problem of designing mechanisms that are both optimal and *individually rational* for hospitals. Under the assumption that donor-patient pairs choose their hospital randomly (according to some distribution), the authors show that any optimal kidney exchange mechanism is likely to be individually rational up to lower-order terms. They also show that, when the hospitals are chosen uniformly at random, with high probability, there exists a matching that is individually rational and optimal up to lower-order terms.

Hamsa Bastani: Analysis of Medicare Pay-for-Performance Contracts.

When health providers are paid for each service they perform as opposed to for the health outcome of their patients, providers are incentivized to prioritize the number of services completed over patient health. To realign the providers' incentives with patient health, Medicare has committed to instead use Pay-for-Performance (P4P) contracts. However, the current P4P mechanisms in use are often designed in an ad-hoc manner. Bastani et al. [2016] investigate mechanisms that achieve optimal utility for Medicare subject to using “small” incentive—an institutional constraint required in practice—and draw conclusions about the various mechanisms that Medicare currently has in place. Among their conclusions, they find that despite initially being more costly for Medicare, offering bonuses instead of just penalizing poor-performers increases Medicare's expected utility. They also find that offering a single performance threshold (above which providers are rewarded and below which providers are penalized) is more effective than a continuous approach, where rewards and penalties vary with the distance from the threshold.

Imanol Arrieta-Ibarra: A Personalized BDM Mechanism for Efficient Market Intervention Experiments.

Measuring the efficacy of market interventions is a corner-stone problem in many social good domains. This is especially the case when implementing a randomized control trial is difficult or even impossible—a setting frequently found in developing world contexts. The BDM mechanism was proposed in the 1960s as a way to both evaluate the benefits of an intervention and measure demand in these situations [Gordon M. Becker, 1964]. Ibarra and Ugander [2017] present a personalized extension of this mechanism, which takes into account settings with heterogeneous consumers. Combining insights from algorithmic game theory and machine learning, they show that this personalized variant can result in a lower cost to the experimenter compared to natural baselines and yield more efficient resource allocation, while keeping statistical soundness.

Yunan Li: Mechanism Design with Financially Constrained Agents and Costly Verification.

Li [2017] considers the setting where willingness to pay (value) and ability to pay (budget) are distinct and private to the buyer. The seller acts like a government agency trying to maximize surplus, and has the ability to verify the reported budget from any agent, but at some cost. This is relevant in settings such as low-income housing allocation or health insurance, where the government aims to mitigate economic inequality by subsidizing costs for those who have less of an ability to pay. Li finds that the surplus-optimal mechanism can be implemented as follows: bidders report their budgets, then the seller provides

cash subsidies based on these reports and assigns the goods randomly. Then, there is a second resale stage, with budget-dependent taxes for reselling. In essence, buyers with lower budgets receive higher subsidies in the first round, but face higher taxes in the second round, and are inspected randomly. This implementation of the optimal mechanism parallels the welfare program used in practice for housing allocation in Singapore.

Peng Shi: Assortment Planning in School Choice. Another seminal problem in mechanism design for social good—school choice—now sees application in various cities across the nation. Systems utilizing school choice mechanisms work by eliciting preferences from students and schools over one another and matching students to schools accordingly. However, these matchings are highly dependent on policy choices made by the school board. These choices (or parameters) include: (1) the number of school options that students are shown, (2) the allowed priorities that schools may have over students, and (3) the quotas for places in the schools. In his work, Shi [2016] reduces the problem of optimizing policy choices to an assortment optimization problem that is a generalization of a common one from the revenue management literature, and he gives new polynomial-time algorithms to solve this problem. Using data from Boston Public Schools, he shows that applying these methods would in fact improve the expected utilities of students and the predictability of the assignment outcome while maintaining the same amount of busing.

Vince Conitzer: Rules for Choosing Societal Trade-offs. When we optimize for societal good, we are required to trade off between different gains and losses which are not directly comparable, e.g. using a gallon of gasoline vs. producing two bags of trash. In addition, when we talk about economic equality, one question that comes up is how to choose our objective function: what *is* social welfare, whose utility matters, and what weight does it get—how do we trade off one person’s utility versus another’s? Conitzer et al. [2016] take a voting-theoretic approach to aggregating how different agents view different trade-offs. They describe properties that an aggregation rule would ideally satisfy, pinpoint that the family of distance-based rules can be interpreted as maximum likelihood estimates, and show that the logarithmic distance-based rule is an especially strong candidate with regard to the properties that a rule should satisfy. They also give nice computational results for arriving at this rule.

Hongyao Ma: Incentivizing Reliability in Demand-Side Response. Mechanism design for energy markets is an under-explored direction with a potential for immense societal impact. There are various incentive issues that arise; recent work has explored how to efficiently decrease energy consumption through proper incentives. Current incentive schemes offer reward-based contracts which often fail to consider consumers heterogeneous consumption profile and reliability. This results in lower participation rates. To address some of this inefficiency, Reshef Meir [2017] propose a VCG-type mechanism for choosing and incentivizing a subset of consumers to reduce consumption in energy markets. Using simple fixed-penalty contracts, the mechanism satisfies incentive-compatibility, results in honest energy reduction preparation efforts, and allows for efficient computation of allocation and prices. They show that it yields higher participation rate, increases reliability, and

reduces expenses compared to current popular mechanisms. The setting is further generalized by Ma et al. [2017], which allows for uncertainty in the agent's costs of responding and multiple levels of preparation efforts. They provide a new, reward-bidding based mechanism under this model and show that this mechanism is also incentive-compatible, reliable, and shows good performance compared to natural baselines.

4. PANEL DISCUSSION

The workshop concluded with a panel discussion showcasing researchers who have extensive experience working at the interface of economics and computation with a social good objective. This panel consisted of Ashish Goel (Stanford University, Management Science and Engineering), Carla P. Gomes (Cornell University, Department of Computer Science), Kevin Leyton-Brown (University of British Columbia, Department of Computer Science), Parag Pathak (Massachusetts Institute of Technology, Department of Economics), and Glen Weyl (Microsoft Research - New England & Yale University). Their work ranges from improving sustainability through algorithmic and optimization techniques to re-envisioning data as labor through radical markets.

Discussions revolved around identifying areas and opportunities where their skill-set provided the most leverage, what unexpected benefits and challenges they have faced working at this interface, what domains provide opportunities for using algorithmic, optimization, and mechanism design techniques to a social good end, and what advice they have for students interested in this field. Many panelists emphasized the need to work closely with domain experts in order to have meaningful real-world impact. Finding a research direction in settings that involve working in interdisciplinary groups can take years, and for many was one of the hardest challenges. Panelists emphasized the need to be flexible and let discussions and experiences guide research directions rather than formulating research questions a priori.

They also discussed initiatives that the economics and computation community can take on, such as a multi-disciplinary review process in order to encourage collaborations across disciplines and subjecting them to evaluation according to the standards of all fields involved, and creating avenues for multi-year projects for researchers at all stages.

5. CONCLUSION

Overall, the workshop was successful in its goals. First, it shared recent work on new social good domains. Second, it introduced the community to new domains and relevant experts. And third, it fostered interest for future work in this area. In particular, by highlighting new research from computer science and related fields side-by-side, the contributed talks demonstrated unexpected breadth in recent social good work.

This body of recent work, the success stories shared on the panel, and the active audience engagement together formed an informative workshop that effectively encouraged the community to contribute to this research agenda.

ACKNOWLEDGMENTS

We are very grateful for the MD4SG'17 Program Committee—Siddhartha Banerjee (Cornell University), Mark Braverman (Princeton University), Nikhil Devanur (Microsoft Research Redmond), Alon Eden (Tel Aviv University), Nicole Immorlica (Microsoft Research - New England), Anna Karlin (University of Washington), Maximilian Kasy (Harvard University), Jon Kleinberg (Cornell University), Irene Lo (Columbia University), Manish Raghavan (Cornell University), Samuel Taggart (Oberlin College), Daniel Waldinger (Massachusetts Institute of Technology), and Matt Weinberg (Princeton University)—for their help with reviewing and chairing sessions at the workshop. We are indebted to EC 2017 Organizing Committee and especially Workshop Co-Chairs Shuchi Chawla (University of Washington, Madison) and Jenn Wortman Vaughan (Microsoft Research-NYC), Local Chair Brendan Lucier (Microsoft Research-New England), and General Chair Costis Daskalakis (Massachusetts Institute of Technology) for all the assistance throughout the proposal, planning, and execution process of the workshop. Finally, we would like to extend our deepest gratitude to our advisors Anna Karlin and Jon Kleinberg for their generosity, wisdom, and support from the very first MD4SG reading group meeting.

REFERENCES

- ABEBE, R. AND GOLDNER, K. 2016. Mechanism Design for Social Good. <http://md4sg.com>. Accessed: 2016-09-16.
- ARNOSTI, N. AND SHI, P. 2017. How (Not) to Allocate Affordable Housing. *Columbia Business School Research Paper*, 17–52.
- BASTANI, H., BAYATI, M., BRAVERMAN, M., GUMMADI, R., AND JOHARI, R. 2016. Analysis of Medicare Pay-for-Performance Contracts.
- BLUM, A., CARAGIANNIS, I., HAGHTALAB, N., PROCACCIA, A. D., PROCACCIA, E. B., AND VAISH, R. 2017. Opting Into Optimal Matchings. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*. SIAM, 2351–2363.
- CONITZER, V., FREEMAN, R., BRILL, M., AND LI, Y. 2016. Rules for Choosing Societal Tradeoffs. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence AAAI Conference on Artificial Intelligence*.
- DELACRÉTAZ, D., KOMINERS, S. D., AND TEYTELBOYM, A. 2016. Refugee Resettlement . Working paper.
- ELLISON, G. AND PATHAK, P. A. 2016. The Efficiency of Race-Neutral Alternatives to Race-Based Affirmative Action: Evidence from Chicago's Exam Schools. Working Paper 22589, National Bureau of Economic Research. September.
- GORDON M. BECKER, MORRIS H. DEGROOT, J. M. 1964. Measuring utility by a single-response sequential method. *Behavioral Science* 9, 226–232.
- HUSSAM, R., RIGOL, N., AND ROTH, B. N. 2017. Targeting High Ability Entrepreneurs Using Community Information: Mechanism Design in the Field . Working paper.
- IBARRA, I. A. AND UGANDER, J. 2017. A Personalized BDM Mechanism for Efficient Market Intervention Experiments. Working paper.

- JONES, W. AND TEYTELBOYM, A. 2016. Refugees' Say. <https://www.refugees-say.com/>. Accessed: 2017-01-19.
- KHO, A. N., CASHY, J. P., JACKSON, K. L., PAH, A. R., GOEL, S., BOEHNKE, J., HUMPHRIES, J. E., KOMINERS, S. D., HOTA, B. N., SIMS, S. A., MALIN, B. A., FRENCH, D. D., WALUNAS, T. L., MELTZER, D. O., KALEBA, E. O., JONES, R. C., AND GALANTER, W. L. 2015. Design and implementation of a privacy preserving electronic health record linkage tool in Chicago. *Journal of the American Medical Informatics Association* 22, 5, 1072–1080.
- LI, Y. 2017. Mechanism Design with Financially Constrained Agents and Costly Verification.
- MA, H., PARKES, D. C., AND ROBU, V. 2017. Generalizing Demand Response Through Reward Bidding. In *Proceedings of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'17)*. 60–68.
- RESHEF MEIR, HONGYAO MA, V. R. 2017. Contract Design for Energy Demand Response. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*. 1202–1208.
- SHI, P. 2016. Assortment Planning in School Choice [Preliminary Draft].
Fischer vs. University of Texas 2016.
Grutter vs Bollinger 2003.
Seattle/Louisville vs. PICS 2007.
- ZHANG, C., ATASU, A., AYER, T., AND TOKTAY, L. B. 2017. Truthful Mechanisms for Medical Surplus Product Allocation.