Design for Weakly Structured Environments

Chapter for *The Future of Economic Design*

Gabriel Carroll, Stanford University*

gdc@stanford.edu

April 18, 2018

Economic theory abounds in simple stylized models. How often we see consumers' preferences parameterized by one-dimensional types; uncertainty represented by two possible states of nature, with agents receiving conditionally independent signals (or perhaps, normally distributed states, observed with normal noise); workers choosing a one-dimensional level of effort.

Economic design, as a field, has aspirations to inform big applications: not only auctions and matching markets, where theory has already been quite influential, but also design of protocols for legislative bargaining, judicial decision-making, or international negotiations; tax codes, patent systems, financial markets. In many cases, practical design in these areas does draw on high-level ideas from economic theory. But if theory is to offer anything like the kind of detail-level contributions it has made in (say) school choice, it will have to confront an obvious gap: Real-world preferences and information are far more complex than in traditional models. Just for an example, in committees or juries that deliberate to make a decision, members do not have a single signal about the likelihood of two states of nature; they may have knowledge ("signals") from various sources militating in favor of one decision or another, but also knowledge about which sources are credible, how to interpret the signals, and so forth. Negotiators may fail to reach instant agreement due not only to private information about the issues they bargain over (and there are typically many simultaneous interacting issues), but also about the

*This essay benefited from discussions and comments from Rohan Pitchford, Kemal Yıldız, and especially Nicole Immorlica. Much of the writing was done while the author was visiting the Research School of Economics at the Australian National University, and their hospitality is gratefully acknowledged. The author is supported by a Sloan Foundation Fellowship.

negotiating process itself and their strength within it. The point is not that we need to add more parameters to existing models; it is that *no* parameterized structural model can hope to capture all such complexities.

It is undoubtedly useful for economists to study stylized models, just as geneticists study fruit flies: as a simple starting point. Structural assumptions make it possible to give closed-form solutions and understand a model exhaustively. But eventually we have to come to the real world, and either understand when and why the simple model provides a good first-order approximation, or identify specific ways in which it is seriously deficient. Either way, we need modeling tools to think seriously about the richness that exists in real design problems, relatively free of structural assumptions.

My expectation — and my hope — is that progress in mechanism design over the coming decades will come from developing more useful general models of preferences, information, and actions in complex environments, relatively free of structural assumptions; and developing conceptual tools to argue for why certain kinds of mechanisms will work well in such environments. An exemplar of such a weakly-structured framework is the Arrow-Debreu model for general equilibrium: the goods can be arbitrary; preferences can take arbitrary form (aside from being monotone, and usually convex). The framework is broad enough to express that some conclusions are quite general (such as the First Welfare Theorem), while others demand much more restrictive assumptions (such as convergence of tatônnement dynamics). Decision theory is another existing area of economics where modeling is usually quite free of structural assumptions; but the typical end point in that field is to give alternative representations or interpretations of a given primitive (choice behavior), and it is not clear whether this methodology can be ported to mechanism design, where the goal is to recommend new mechanisms.

Making this progress toward more free-form models will require a shift in the criteria by which research in economic theory is evaluated. There is a strong cultural norm in economics that favors models that can be solved exactly. This is true both in the axiomatic branch of the design literature — where the ideal is a list of criteria that pick out a unique mechanism, or a small class of mechanisms, for some application — and in the Bayesian branch — where one writes down a numerical objective, typically expected revenue or expected welfare, and solves for the exact optimum. If a theorist proposes a non-optimal mechanism, the instinctive reaction is "plenty of things are non-optimal; what's so special about yours?"

But in any situation that remotely approaches the complexity of most real-world design problems, there is no hope of describing the exact optimum. Instead, design is and should be pluralistic. Economic design is not so different from design of physical products: a consumer may not find a perfect car to drive, or even a perfect sandwich for lunch, but having a range of options on the market makes it more likely she will find a choice she can be pretty happy with; and the same can be said for voting rules or negotiation protocols.

Accordingly, the criterion to evaluate a theoretical contribution should be: does it offer a new idea that will plausibly be helpful in eventually informing design (or in understanding mechanisms that are in use)? This is of course a subjective question. Perhaps the field focuses on exact optimization simply because it is unambiguous whether a particular work meets this criterion. A challenge for the mechanism design theory community, then, will be to reach some amount of agreement on what other kinds of ideas to recognize as valuable contributions, so that future researchers are less compelled to focus on the confining class of problems that have exact solutions.

What might such contributions look like? I will suggest a few pointers for inspiration based on existing work (of course, citations are selective due to space constraints). This is not meant to be an exhaustive list, and I hope that future work will discover new directions.

I find it helpful to classify work in economic design into two main kinds, "principles" and "engineering." "Principles" research uses stylized models to offer general insights. For example, the Myerson-Satterthwaite theorem [19] showed us that there is generally no way to ensure efficient outcomes when parties have private information about preferences; subsequent work such as [10] and [22] showed how this conclusion can be overturned if the outside option is sufficiently symmetric. "Engineering" works out technical details to tailor mechanisms to specific contexts (which can mean either designing mechanisms for specific real-world uses in practice, or applied-theory models inspired by an application). Thus, for example, the two-sided matching literature, which has produced a stream of research adapting versions of the Gale-Shapley algorithm to the idiosyncrasies of different applications (examples include [1, 13, 23]).

Both of these branches can contribute to designing for weakly structured environments. On the principles side:

- One kind of contribution will consist of offering new tools for the designer's toolkit. An example is Chassang's work on calibrated contracts [7]. He considers a very general dynamic moral hazard model, with almost no structural assumptions. The principal would like to incentivize an agent (a wealth manager) to make good investment decisions, using a linear contract, so that the agent shares in the gains and losses he produces. But limited liability prevents her from making the agent share the losses. Chassang adapts tools from no-regret learning to show how this infeasible linear contract can nonetheless be closely approximated by adjusting the share fraction up or down over time depending on past performance.

Thus, the emphasis is not on finding an optimal contract, but on providing a tool that can plausibly be useful for getting around a specific obstacle in incentive design.

- There are a number of approximation results that show how simple mechanisms can perform reasonably well across a broad range of environments. For just one of many examples, consider selling two goods to a buyer with unknown values. Finding the revenue-maximizing mechanism is a difficult and ill-understood problem in general; unlike with just one good, the optimum for two goods can be quite complicated, e.g. offering infinitely many different probabilistic bundles at various prices [11]. But Hart and Nisan [14] showed that if the values for the goods are independently distributed, then just selling the goods separately gives at least 1/2 the maximum possible revenue. Subsequent works have extended this in many directions (for example [4, 21]).

  Related to this is the price-of-anarchy literature, studying welfare in equilibria of simple mechanisms with multiple agents. For example, [8] showed that when m buyers buy n goods via simultaneous second-price auctions held separately for each good, welfare in Bayes-Nash equilibrium is a 1/2-approximation of first-best welfare as long as buyers' preferences over goods are submodular. This, too, has spawned many follow-ups and generalizations ([20] gives a survey).

  Such approximation results are bread and butter in computer science. In economics they are often greeted with the skeptical reaction that getting 1/2 of the optimum is not very satisfying. I might point out in reply that when exact optimization is hard and the starting point is zero, getting to 1/2 is a pretty decent first step. In any case, the usual criterion for evaluating economic theory is not whether the results map literally into practice — they rarely do — but whether the analysis gives new (and perhaps even surprising) ideas that help us better understand the situation at hand. By this criterion, approximation results surely have a role to play. As Hartline [15] argues, by studying whether a particular class of mechanisms can or cannot give an approximation guarantee, we have a principled way to formalize ideas about which properties of mechanisms are of first-order importance and which ones are not. In addition, the technical analysis underlying approximation guarantees sometimes leads to new conceptual frameworks for understanding classes of mechanisms, such as the "smoothness" framework for price-of-anarchy results.

- Some headway can also be made by considering models that retain strong structure on some dimensions while dropping it on others. An example of this is the work of Bergemann, Brooks, and Morris [3]. They study a traditional first-price auction for a single object, but make no structural assumptions about the players' information (either about their own value for the object, others' values, or others' information). They show how one can nonetheless obtain lower bounds for equilibrium revenue in such a mechanism. Intuitively, the revenue cannot be very low, since then any bidder could simply deviate to a moderately low bid, and win at a profit with high probability. It turns out that this idea can be pushed to its logical conclusion to give a tight lower bound for revenue as a function of the value distribution, and to describe the worst-case information structure. The fact that a tight bound can be obtained seems specific to the application, but the idea of allowing information to be unstructured is in line with the theme here. ([6] performs an exercise in a similar spirit with a binary decision mechanism.)

On the engineering side, it is hard to envision details about specific applications at a distance. Yet, there may be some categories of methodological advances that are likely to have widespread impact across many applications. I will make a few guesses:

- For mechanisms that are used repeatedly, it is natural to try to learn about distributions over agents' preferences, information, and so forth based on previous runs of the mechanism, and dynamically adjust parameters of the mechanism accordingly. This creates a need for tools to efficiently learn these parameters from past data; and indeed, there is demand for such tools in e-commerce. Recent theoretical work on data-driven design, such as [9, 18], explores this topic in the canonical auction model.

  One could readily envision expanding these ideas to other domains of incentive design where there are extensive past interactions to learn from, such as government procurement [16] or regulation of utilities [2]. In these domains, more so than in e-commerce auctions, interactions are complex, and the prevailing theoretical tools are stylized models that are hard to interpret literally. Hence, many new questions arise about what a designer can learn from quantitative data about the structure of the environment, as well as how best to learn it.

- In some problems, just finding a good outcome is already an engineering challenge. This challenge by itself falls within the traditional territory of operations researchers, not economists. Yet, there are situations where incentives interact with the optimization problem in substantive ways, creating new theoretical challenges, for which economics has something to contribute. This has been a theme of theoretical research in algorithmic game theory for several years, but is just starting to be a serious consideration in practical applications, with the recent FCC incentive auction to reallocate radio spectrum in the United States as a headline example [17]. As both the theory of mechanism design and the adoption of algorithms to automate decision-making move into more complex domains, the interaction between optimization and incentives is likely to become more widespread.

- The design of the interface by which participants interact with a mechanism is also important. Currently, much of mechanism design starts by applying the revelation principle and assumes that agents' full preferences can be elicited. As design moves toward applications in which preferences and information are more complex and unstructured, this assumption will be less and less realistic, and it will be increasingly useful to have principled ways of thinking about how to efficiently elicit the most important information.

  Also important, and currently more neglected, is communication in the other direction: especially if one hopes to engineer automated decision-making mechanisms for settings where participants can refuse to accept the mechanism's output (or can stop using it for future interactions), one challenge can be explaining clearly to participants why they can't have everything they want, and persuading them to be satisfied with the outcome of the mechanism. As technology develops to bring increasingly detailed and quantitative input into large-scale collective decisions (such as participatory budgeting, [5, 12]), this latter task may well also become technical enough to have a role for formal theory.

Finally, given that I have started from an ambitious vision of the scope of possible applications for mechanism design, I feel it is important to temper ambition with humility: Some of the areas of application I have described are relatively removed from the traditional domain of economics — the allocation of material resources — and are more the province of experts in political science, law, or social psychology. While I have predicted that economics will, by developing more general and flexible models, increase its ability to inform actual design in these areas, economics will not and should not supplant domain-specific expertise. Economics has some specific comparative advantages to offer: the ability to think quantitatively about cost-benefit tradeoffs; the habit of thinking about strategic incentives, and especially about how changing the rules of the game will make the players play differently. These strengths should complement the perspectives and tools that come from other disciplines.

References

[1] Itai Ashlagi, Mark Braverman, and Avinatan Hassidim (2014), "Stability in Large Matching Markets with Complementarities," *Operations Research* 62 (4): 713–732.

[2] David P. Baron and Roger B. Myerson (1982), "Regulating a Monopolist with Unknown Costs," *Econometrica* 50 (4): 911–930.

[3] Dirk Bergemann, Benjamin Brooks, and Stephen Morris (2017), "First-Price Auctions with General Information Structures: Implications for Bidding and Revenue," *Econometrica* 85 (1): 107–143.

[4] Moshe Babaioff, Nicole Immorlica, Brendan Lucier, and S. Matthew Weinberg (2014), "A Simple and Approximately Optimal Mechanism for an Additive Buyer," *FOCS* '14: 21–30.

[5] Yves Cabannes (2004), "Participatory Budgeting: A Significant Contribution to Participatory Democracy," *Environment and Urbanization* 16 (1): 27–46.

[6] Gabriel Carroll (2016), "Informationally Robust Trade and Limits to Contagion," *Journal of Economic Theory* 166: 334–361.

[7] Sylvain Chassang (2013), "Calibrated Incentive Contracts," *Econometrica* 81 (5): 1935–1971.

[8] George Christodoulou, Annamária Kovács, and Michael Schapira (2016), "Bayesian Combinatorial Auctions," *Journal of the ACM* 63 (2): #11.

[9] Richard Cole and Tim Roughgarden (2014), "The Sample Complexity of Revenue Maximization," *STOC* '14: 243–252.

[10] Peter Cramton, Roberts Gibbons, and Paul Klemperer (1987), "Dissolving a Partnership Efficiently," *Econometrica* 55 (3): 615–632.

[11] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos (2013), "Mechanism Design via Optimal Transport," *EC* '13: 269–286.

[12] Ashish Goel, Anilesh K. Krishnaswamy, Sukolsak Sakshuwong, and Tanja Aitamurto (2016), "Knapsack Voting: Voting Mechanisms for Participatory Budgeting," unpublished paper, Stanford University.

[13] Isa E. Hafalir, M. Bumin Yenmez, and Muhammed A. Yildirim (2013), "Effective Affirmative Action in School Choice," *Theoretical Economics* 8 (2): 325–363.

[14] Sergiu Hart and Noam Nisan (2017), "Approximate Revenue Maximization with Multiple Items," *Journal of Economic Theory* 172: 313–347.

[15] Jason Hartline (2012), "Approximation in Mechanism Design," *American Economic Review* 102 (3): 330–336.

[16] Jean-Jacques Laffont and Jean Tirole (1986), "Using Cost Observation to Regulate Firms," *Journal of Political Economy* 94 (3): 614–641.

[17] Kevin Leyton-Brown, Paul Milgrom, and Ilya Segal (2017), "Economics and Computer Science of a Radio Spectrum Reallocation," *Proceedings of the National Academy of Sciences* 114 (28): 7202–7209.

[18] Jamie H. Morgenstern and Tim Roughgarden (2015), "On the Pseudo-Dimension of Nearly Optimal Auctions," *NIPS* '15: 136–144.

[19] Roger B. Myerson and Mark A. Satterthwaite (1983), "Efficient Mechanisms for Bilateral Trading," *Journal of Economic Theory* 29 (2): 265–281.

[20] Tim Roughgarden, Vasilis Syrgkanis, and Éva Tardos (2017), "The Price of Anarchy in Auctions," *Journal of Artificial Intelligence Research* 59: 59–101.

[21] Aviad Rubinstein and S. Matthew Weinberg (2015), "Simple Mechanisms for a Subadditive Buyer and Applications to Revenue Monotonicity," *EC* '15: 377–394.

[22] Ilya Segal and Michael D. Whinston (2011), "A Simple Status Quo that Ensures Participation (with Application to Efficient Bargaining)," *Theoretical Economics* 6 (1): 109–125.

[23] Tayfun Sönmez and Tobias B. Switzer (2013), "Matching with (Branch-of-Choice) Contracts at the United States Military Academy," *Econometrica* 81 (2): 451–488.