Research Statement

Ruggiero Cavallo
Computer Science
Harvard University
cavallo@eecs.harvard.edu

Deriving effective group decision-making procedures for complex and potentially uncertain environments is hard, but of fundamental importance, and the challenges grow significantly more daunting when individuals in a group are *self-interested*. There is an inherent tension in striving to achieve *social* goals in decisions that will impact individuals that are only concerned with *local* objectives. Innumerable scenarios fit this mold, from resource allocation to coordinating behavior in the presence of global constraints: a government seeks to allocate wireless spectrum to the company that will derive greatest value from it; an administrator seeks to coordinate the activities of lazy subordinates towards company efficiency, etc. Solving such problems is the central focus of my research.

Typically the metrics used to evaluate the quality of any given decision that will impact a group inherently depend on information that is private to the individual parties. A classic result shows that, in order to reach decisions that satisfy desirable social objectives, one must be able to elicit truthful reporting of this private information by the agents, i.e., disarm the individuals' incentive to strategically manipulate the system. This can often be achieved via carefully defined monetary transfer payments imposed on individuals; however, in many cases determining the appropriate payments, or even evaluating prospective decisions, is *computationally* challenging.

Thus this area—computational mechanism design—is right on the border of economics and computer science. My dissertation research attacks two fundamental classes of decision-making problems at this intersection:

1. How can the behavior of a group of agents in a *dynamic* and uncertain world be organized to achieve good system-wide outcomes, despite agents acting only to maximize their own individual welfare? In the real world, decisions do not exist in an isolated instant, but rather are situated in a temporal context with other decisions. Individuals will act to maximize their *utility over time*, and decisions in the present influence what the world will look like in the future, but rarely in completely predictable or deterministic ways. The dual challenges of computing optimal decision policies and incentivizing agents to continuously participate honestly must be dealt with simultaneously.

2. What decision-making mechanism can be implemented to maximize the welfare a group of self-interested agents jointly derive from the outcome? The payments mechanism design typically requires—to elicit truthful revelation of private information—are made *from* the agents *to* some central coordinator; but from the agents' perspective, such payments are waste, and it is natural to seek to minimize them. It was previously presumed by many that the ubiquitous VCG mechanism is the best one can do; my work provides a mechanism applicable to arbitrary group decision-making problems that shows otherwise.

Looking ahead, many significant challenges remain in moving from theory to practice. To name just one, *optimal* decision policies are typically the only ones that can be implemented in equilibrium, yet, in many domains, computing them is intractable. I am developing a theory of how to do mechanism design in such intractable environments. In tandem, I am excited about making mechanism design work to involve agents more deeply in its execution, yielding systems in which the stakeholders are invested and involved in the coordinated computation of an optimal decision policy and appropriate transfer payments, despite their inherent self-interest.

Current work

1. Mechanism design for dynamic environments

The field of mechanism design (MD) addresses decision-making problems involving self-interested parties (or "agents"), in which the goal is to reach an outcome that meets certain system-wide criteria. The crux of the challenge is agent self-interest, combined with the fact that—almost always—evaluating an outcome's utility to an agent requires knowledge of some information privately held by that agent. A *mechanism* consists of a decision function and a payment function, and solutions are achieved by defining these functions in a way that aligns the incentives of the agents towards the desired outcome. For instance, one may wish to allocate a single, indivisible good to an individual (in a group) that would benefit most from obtaining it; to do so requires knowledge about the agents' valuations. One solution is the second-price or "Vickrey" auction, in which each agent is always best-off reporting (bidding) its true valuation to the auctioneer (or "center"); determining the optimal allocation is, then, a simple matter of comparing the bids.

In my dissertation I extend the mechanism design enterprise to *dynamic* settings, those in which a sequence of decisions is to be made over time. As an example, imagine a natural extension of the scenario described above in which a single resource is to be allocated *repeatedly*—an agent is granted the resource, uses it for a time, returns it to the center, and this cycle repeats. Agent valuations may *change over time*; in particular, the agent that obtains the resource in any given epoch may learn something about how much value the resource bears, or perhaps the good is useful to the agent only for a limited number of epochs.

Previous work in "online mechanism design" added important but limited dynamics to the standard, static mechanism design setting: a changing population of agents is considered, where each agent obtains no new private information after arrival. In many (or most) practical settings, though, as in the example above, agents *will* acquire new private information throughout execution of the mechanism. In [4] I presented the first socially-optimal mechanism that provides incentives for agents to report this private information as they acquire it, whenever they acquire it, in dynamic settings. This work, which is collaborative with my advisor David Parkes and Satinder Singh (University of Michigan), extends the intuition that brought successes in static mechanism design—if the payoffs of the agents can be redefined (via transfer payments) such that each agent is best off when the system as a whole is best off, agents will participate honestly. The solution I proposed defines payments such that each agent's expected utility in equilibrium from the start of the mechanism equals its expected contribution to social welfare. The mechanism also has the property of being budget-balanced in expectation (i.e., the expected net transfers between the center and the agents are 0)—however, the price of achieving this property is a weaker participation (individual rationality) property, a significant drawback. In [5] I presented an application of these ideas to a semi-autonomous robot coordination problem.

Subsequent to publication of [4], in an as yet unpublished paper Bergemann & Välimäki[1] proposed a mechanism with equally strong incentive properties and a stronger participation property—the mechanism can be considered the natural analogue of the VCG mechanism (which reduces to a Vickrey auction in single-item allocation problems) for a dynamic setting, and is thus termed *dynamic-VCG*. Dynamic-VCG constitutes a very satisfactory solution to the most basic problem, where, for instance, there is an auctioneer allocating goods and a fixed population of agents. But many important questions remain. I proposed a mechanism that extends dynamic-VCG to handle a changing population of agents, and scenarios in which agents are periodically

[1]Dirk Bergemann and Juuso Välimäki. Efficient dynamic auctions. *Cowles Foundation Discussion Paper 1584, http://cowles.econ.yale.edu/P/cd/d15b/d1584.pdf*, 2006.

"inaccessible" [6]. I believe this work, also in collaboration with Parkes and Singh, represents significant progress towards making dynamic mechanisms more applicable to real problems, as the proposed mechanism manages, for instance, cases in which communication between agents and the center is faulty (a communication breakdown leads to intermittent inaccessibility). Interestingly, this work also unifies the complementary theories of online-MD and dynamic-MD under one framework.

2. Redistribution mechanisms

The usual objective pursued in mechanism design is maximization of social welfare (also termed efficiency), i.e., the sum of agent utilities generated by a decision. Somewhat ironically, typical MD solutions often require agents to make payments to the center so large that most of the utility yielded from the decision is lost to the agents. I addressed the question of whether we can do better in [2], and answered it affirmatively for cases in which we can reason explicitly about the context of the decision to be made, e.g., by taking into account the fact that a decision may necessarily yield utility for any one individual at the expense of others.

In order to make application of a mechanism practical, budget and participation requirements are typically imposed—the mechanism should not run a deficit, i.e., net monetary transfers should not flow from the center to the agents (no-deficit), and agents should be guaranteed that they won't achieve negative utility from participating in the mechanism (ex post individual rationality, or IR). The VCG mechanism—in which each agent must pay the center an amount equal to the negative externality it imposes on other agents—implements an efficient decision policy in dominant strategy equilibrium, and at the same time satisfies both the no-deficit and ex post IR constraints. However, in many cases VCG is not good at all for the agents. For instance, consider a single-item allocation problem in which there are 4 agents, with values 7, 8, 9, and 10, respectively, for the item. The VCG mechanism (the Vickrey auction, here) will allocate the item to the agent with value 10, but requires that he pay the center 9; 90% of the utility from allocation goes to the center, or is "wasted" from the perspective of the agents.

In [2] I examined the problem of minimizing this waste within the context of an efficient, no-deficit, ex post IR mechanism, thus maximizing the welfare of the individuals that seek to reach a decision, while at the same time maintaining all the desirable properties of the VCG mechanism. What I discovered was somewhat unexpected—it turns out that the VCG mechanism is the best one can do when no "domain information" is taken into account, but when such information is considered one can do significantly better. For instance, in single-item allocation problems there is significant structure to agent valuations that can be exploited; specifically, it is typically know—regardless of what the agents bid—that any agent who is not allocated the item will obtain no utility from the decision. This observation alone allows for implementation of a mechanism in dominant strategy equilibrium that keeps a far greater portion of the utility within the group of agents. Considering again the 4-agent, single-item allocation example above, the mechanism I proposed maintains a utility of 8.5 (85%) within the group of agents, while VCG maintains a utility of only 1 (10%).

Beyond simply doing better than previously known mechanisms, I showed that my proposal is strongly optimal when a particular fairness (or "anonymity") condition is imposed. Among all mechanisms that satisfy this fairness condition, none *ever* (i.e., for any set of agent preferences over outcomes) yields higher payoff to the group of agents. I also explained the relationship between this fairness condition and other more standard ones that may be required in its place, and demonstrated that when no fairness condition is imposed there is no mechanism that is optimal in the strong sense I just described.

3. Mechanism design for intractable environments

I am currently developing a framework for analyzing and managing incentives in environments in which perfect solutions are not available, either to the coordinator or the agents [1]. Such environments abound. For instance, in combinatorial allocation problems[2] the complexity of determining an optimal allocation is exponential in the number of goods to be allocated. While state of the art solvers can handle large problem instances, clearly there are many cases in which exact solutions are unavailable. To take another pertinent example close to my heart, exact solutions are also typically unavailable in dynamic environments, in which computing an optimal decision problem amounts to solving a multi-agent Markov decision process, where the state space is exponential in the number of participating agents. This work is an attempt to bridge a large part of this gulf that separates the theory of mechanism design from practice.

The intuition underlying my approach is this: though implementing an optimal decision policy may be intractable, computing payments that align the interests of all agents towards achieving that goal is *always* tractable. Then, given such a payment scheme, it can be shown that agents will deviate from truthful participation only if they know of a decision procedure that will yield more social welfare than the one chosen by the coordinator. I propose a new style of decision procedure, called *belief coordination mechanisms*, in which agents can share knowledge of heuristic decision policies with the coordinator. This process yields, in an ex post equilibrium, implementation of a policy that is at least as good as any heuristic known to the coordinator.

4. A combinatorial exchange, and a language for describing preferences

In a massive project somewhat distinct from my dissertation research, I was part of significant collaboration that designed and implemented the first fully expressive *iterative combinatorial exchange* (ICE) [7]. A combinatorial exchange generalizes a combinatorial auction to the case where there are multiple buyers and sellers, and agents who are both buyers and sellers. This generalization is significant: a combinatorial exchange allows for *trades*, where an agent may have preferences of the type "I will sell good A if and only if I am allocated goods B and C". Our system is *iterative*, meaning it proceeds in a series of rounds; in each round agents are asked to reveal a limited amount of information about their preferences over the goods, and then a provisional allocation and prices are determined. In the following round this provisional information guides agents in revealing further information (e.g., at these prices, I prefer an alternate bundle of goods to the one I have been provisionally allocated). The iterative nature of the system, combined with carefully designed queries of agents, allows agents to keep private the information that is not relevant to determining the efficient allocation; it is also key for managing large problems in which it may be impossible for the agents themselves to precisely know their preferences over all possible allocations.

Arguably the most significant innovation of the system—and the primary area of my contribution—is the *preference representation language* that was designed for the exchange. We developed a tree-based bidding language (*TBBL*) [3] that has many novel features. Bidders' preferences are represented in a tree structure, where the leaves represent atomic goods and internal nodes correspond to operators for combining subsets of goods. The semantics of the language allows for very concise representation of very complex preferences. For instance, the language provides a general operator for internal nodes that allows one to express things like "I have value between 2 and 4 for good A, and I will get an additional value of between 5 and 10 if I acquire at least two

[2] In a combinatorial allocation problem, agents have preferences over *bundles* of goods; so, for example, if there are 3 goods, A, B, and C, agents may have arbitrary preferences for all 7 distinct bundles: A, B, C, AB, AC, BC, and ABC.

goods in set $\{B, C, D\}$". In [3] I designed an updated version of TBBL with basic extensions, and showed that it subsumes the other main languages that have been proposed for combinatorial allocation problems, i.e., it can represent *any* preferences at least as concisely.

Future work

I view the work in my dissertation as progress towards making the theory of mechanism design more effective and applicable in the real world, even though the results are—in the main—fundamental, abstract, and theoretical. The real world is not static, and so a static theory of decision-making with self-interested agents can only go so far. In the real world it is also the case that, frequently, agents will not be satisfied by "solutions" that do not principally benefit the group of interested parties; in this regard there is strong reason to consider my redistribution mechanisms more plausibly adoptable than other proposals. In the future I want to continue this enterprise—my forte is in theory, but I want to work to develop theories that make progress for problems actually being faced in reality.

1. When standard assumptions break down

Computation (with a note about multi-armed bandits)

Perhaps the most important assumption lurking in standard mechanism design theory is that a socially optimal decision can be computed. As I discussed earlier, this assumption is invalid in most realistic dynamic settings, and also in important static settings. My work on belief coordination mechanisms [1] establishes a framework for approaching such problems in the static case, but extending the theory to dynamic settings will be essential. One exciting direction is the prospect of applying belief coordination mechanisms to the multi-agent learning, planning, and acting techniques that come out of traditional AI research. Designing a dynamic mechanism in which agents learn and act in accordance with, for instance, a Q-learning style heuristic algorithm would constitute a significant advance for the theory of decision-making in complex, uncertain environments with self-interest.

At the same time, work should be done to discover optimal solutions where they are available, and to identify real-world scenarios in which optimal solutions can be applied. For instance, the work of Gittins[3] provided an efficient solution to the multi-armed bandit learning/decision-making problem—this result, in turn, implies the optimality and implementability of a dynamic mechanism for repeated allocation of a single item. The import of Gittins's achievement is hard to overstate, yet still, there are many significant closely related questions that remain unresolved. For instance, the Gittins result only holds when agents "discount" the future, i.e., value future reward less than reward received in the present. There are important economic problems that fall in an undiscounted setting, and a more complete theory of optimal decision-making is needed there.

[3] J. C. Gittins and D. M. Jones. A dynamic allocation index for the sequential design of experiments. In *Progress in Statistics*, pages 241–266. J. Gani et al., 1974.

Utility functions and budget constraints

Practically omnipresent in work in computational mechanism design are assumptions that agents have quasi-linear utility functions (i.e., each agent's utility is basically some intrinsic value for an outcome plus whatever money it is payed), and that agents are not budget-constrained (i.e., there is no price that an agent would be incapable of paying). Neither of these assumptions are completely realistic, and I want to make the theory of mechanism design more applicable when they break down.

2. Distributing computation and decision-making

A multi-agent setting in which a complex problem must be solved naturally leads to questions of distributing computation. When self-interest is involved, things get tricky, as agents may have incentive to shirk computing responsibilities or misreport the results of computation, and privacy concerns must be considered. In some problem domains independence relations mean policy computation is inherently factored, making the prospect of distributed solutions more hopeful. For instance, looking again at the repeated allocation of a single good, the solution I proposed in [4] incentivizes agents to both honestly compute their own "Gittins index"[4] and report it to the coordinator. The center's job is reduced to comparing these indices, allocating the good to the agent with highest reported index, and computing transfer payments.

In ongoing work in collaboration with David Parkes, I am extending this idea to establish a mechanism in which agents will honestly perform and report the results of costly "deliberation" about their valuations; when agents can research new uses (with uncertain results) for a resource to be allocated, the optimal policy will often have them do so (to an extent depending on other agents' valuations) prior to choosing an agent for allocation. We provide a solution in which self-interested agents will choose to do deliberation when and only when it is prescribed by the socially optimal policy, and will report the true results of such deliberation so that an optimal allocation decision can be made. However, the methods I employ in that work do not apply to what is perhaps an even more compelling problem: if agents can *learn better estimates*[5] of their valuation of a good (e.g., via some sampling process), can we specify a computationally efficient mechanism that incentivizes them to do so in a way that is *socially* optimal? A solution to this problem seems likely to require an advance in the theory of optimal control of independent Markov processes, and perhaps further advances on the mechanism design side as well.

Questions like this are part of what can be considered a broader program that seeks to apply mechanism design beyond the standard realm of eliciting a truthful valuation report—the incentives must go deeper, and distributed computational problems must be solved along the way.

3. Fairness

In order to make mechanism design more plausible and applicable in the real world, I believe questions of fairness and equitability must be given greater attention. The work I've done

[4]Gittins indices are values of a function described by Gittins that can be used to determine an efficiently computable and optimal decision policy: choose the action with highest Gittins index.

[5]This is distinct from the deliberation setting, where agents have a value certain for a resource and a model (with uncertainty) of how their value might *increase* from performing research/deliberation; in the learning setting, agents are unsure what value would be yielded from the resource, and would not, independent of other agents, derive any benefit from learning a better estimate before receiving the good.

on redistribution mechanisms can be seen as partially in this vein, as it ameliorates a major drawback to the typical MD solution, where the majority of the utility from a decision is often "unfairly" transferred away from the agents.

There is a significant literature on the fair division of goods, with many possibility and impossibility results on the simultaneous satisfaction of criteria like economic efficiency, equitability, and envy-freeness. Interestingly, however, this literature rests on a set of assumptions, particularly regarding the utility functions of agents, that are different than those in classic mechanism design, and it does not typically consider the possibility of transfer functions as a solution technique. Assumptions here should really not be considered weaknesses of a theory, as the solutions inherently depend on the properties of the system under analysis (e.g., the *actual* utility functions of the participating agents). But I would like to explore the interface between these two literatures, and I hope to find solutions that are more applicable, fair, and beneficial to the stakeholders in real world problems.

References

[1] Ruggiero Cavallo. Belief coordination mechanisms: Group decision-making procedures for intractable environments. Working paper, http://eecs.harvard.edu/econcs, 2006.

[2] Ruggiero Cavallo. Optimal decision-making with minimal waste: Strategyproof redistribution of VCG payments. In *Procceedings of the 5th International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS'06), Hakodate, Japan*, 2006.
Nominated for best student paper award.

[3] Ruggiero Cavallo, David C. Parkes, Adam Juda, Adam Kirsch, Alex Kulesza, Sebastien Lahaie, Benjamin Lubin, Loizos Michael, and Jeffrey Shneidman. TBBL: A tree-based bidding language for iterative combinatorial exchanges. In *IJCAI-05 Multidisciplinary Workshop on Advances in Preference Handling, Edinburgh, Scotland*, 2005.

[4] Ruggiero Cavallo, David C. Parkes, and Satinder Singh. Optimal coordinated planning amongst self-interested agents with private state. In *Proceedings of the Twenty-second Annual Conference on Uncertainty in Artificial Intelligence (UAI'06)*, pages 55–62, 2006.

[5] Ruggiero Cavallo, David C. Parkes, and Satinder Singh. Optimal coordination of loosely-coupled self-interested robots. In *the Workshop on Auction Mechanisms for Robot Coordination, AAAI-06, Boston, MA*, 2006.

[6] Ruggiero Cavallo, David C. Parkes, and Satinder Singh. Online mechanisms for persistent, periodically inaccessible self-interested agents. Working paper, http://www.eecs.harvard.edu/econcs, presented at the *DIMACS Workshop on the Boundary between Economic Theory and Computer Science*, October 2007.

[7] David C. Parkes, Ruggiero Cavallo, Nick Elprin, Adam Juda, Sebastien Lahaie, Benjamin Lubin, Loizos Michael, Jeffrey Shneidman, and Hassan Sultan. ICE: An iterative combinatorial exchange. In *Proceedings of the Twenty-second Annual Conference on Uncertainty in Artificial Intelligence (UAI'06)*, pages 249–258, 2005.

An expanded version of this paper, by Lubin, Juda, Cavallo, Lahaie, Shneidman, and Parkes, has been accepted to JAIR subject to revisions.