

## Research Statement

**Ruggiero Cavallo**

Microsoft Research, New York City

cavallo@post.harvard.edu

December 2012

Reaching decisions that serve the common good—for small communities, companies, or even whole societies—is a goal that all may embrace, yet in practice is often stymied by the self-interest of the individuals concerned and the computational complexity of optimal choice. My research passion is to design tools to help groups overcome this disconnect, to engineer rules of interaction that serve our social-welfare driven aspirations by disarming the influence of selfishness, bringing inclinations to think and act locally into harmony with tractable higher-order systems of global optimization.

Innumerable scenarios manifest this tension between individual and group interests, from resource allocation to coordinating behavior in the presence of global constraints: a government seeks to allocate wireless spectrum to the companies that will derive greatest value from it; an administrator seeks to coordinate the activities of subordinates towards company efficiency; an online “crowd-sourcing” labor-marketplace designer seeks to efficiently connect producers and consumers, etc. A powerful tool for aligning individual interests is monetary transfer payments, as in an auction; however, in many cases determining appropriate payments, or even evaluating prospective decisions, is *computationally* challenging. Moreover, important second-order factors such as budget-balance, fairness, and the distribution of welfare throughout the population must be addressed.

Many critical problems in this area, called *mechanism design*, lie on the border of computer science and economics. My research is centered here and has dimensions that fall under *multi-agent systems*, *electronic commerce*, and *artificial intelligence*. The area is vibrant: fundamental questions remain unanswered, while new applications of increasing scope and importance are consistently arising, particularly with the growth of the internet and online commerce. On the theory side, the following are a few of the fundamental mechanism design questions I have addressed:

- Can we reach welfare-optimal decisions without imposing large transfer payments on the agents? The payments classic mechanisms require of individuals—to incentivize behavior conducive to group welfare—are welfare-*diminishing* in their own right. How can we minimize them? (Consider government allocation of public resources, where civic welfare is the goal.)
- How can the behavior of a group of agents in a *dynamic* and uncertain world be organized to maximize social welfare? The dual challenges of computing optimal decision policies and incentivizing agents to continuously implement them must be addressed simultaneously. (Consider a multi-agent Markov decision process with *selfish* agents.)
- When a group of agents in an uncertain environment share a common goal but disagree about the best way to accomplish it, can incentives be provided to achieve a decision policy that is optimal according to an aggregate of their beliefs? (Consider a CEO that must choose among a set of business proposals, and seeks honest opinions from company VPs.)

Mechanism design needs to become more useful through adoption of more *realistic models*. One branch of my research addresses empirically-motivated deviations from classic theory, incorporating factors such as altruism and subjective belief formation. Grappling with the realities of human behavior is essential to making mechanism design theory broadly relevant and effective, and this project is becoming more and more feasible as access to behavioral data explodes.

My research also has a more specifically applied dimension focusing on emerging internet applications, including crowdsourcing procurement markets and auctions for online advertising. Agent-mediated commerce is opening the door to unprecedented opportunities for large-scale deployment of AI and multi-agent systems research informed by game theory, and here again, there is a key empirical facet driven by the vast data becoming available. I have obtained insight into these issues through my recent experience at two of the world’s most prominent tech-industry research labs.

## I Some major research themes

Mechanism design can be viewed as the engineering side of game theory: for agents that are selfish and reason about each other’s behavior, how can we design a set of rules that leads to optimal group outcomes? In a decision problem, the crux of the challenge is the combination of self-interest and private information—evaluating an outcome’s utility to any given agent requires information that the decision maker doesn’t have. We may ask agents to share this information, but will they do so truthfully? A *mechanism* specifies a decision function (a way of choosing outcomes) and a transfer function (a way of defining incentive payments to be made to or from agents) that take as input the claims agents make about their private information; truthful reporting is 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 the person in a group that would benefit most from obtaining it; to do so requires determining the agents’ values. In the second-price or Vickrey auction solution to this problem, each agent is always best-off “bidding” its true valuation. Determining an optimal allocation is then a simple matter of comparing the bids.

While traditionally a microeconomics discipline, mechanism design has found a very natural place in the research agendas of a growing community of computer scientists. Part of this phenomenon has to do with the *engineering* nature of the task; part with the algorithmic dimension of determining optimal decision policies, computing appropriate incentives, and demonstrating that they satisfy the domain’s constraints; and also a large part with the growing complexity of the application areas—with significant roles for learning and other aspects of AI—that have especially blossomed online.

### Redistribution mechanisms

One of the primary mechanism design objectives is maximization of *social welfare*, i.e., utility to the agents. However, classic solutions such as the Vickrey–Clarke–Groves (VCG) mechanism often require agents to make payments so large that most of the value is lost to them. In [2] I addressed the question of whether we can do better, and answered it affirmatively for settings in which we can reason explicitly about the context of the decision to be made, for instance in resource allocation problems, where agents that receive no goods are known to obtain zero value. Such structure allows us to determine large lower-bounds—*independent of any one agent*—on the amount of revenue that VCG would yield, which in turn allows us to “redistribute” revenue back to the agents without distorting incentives. In single-item allocation, whereas the Vickrey auction requires the second-highest value to be paid by the high-bidder (an amount that one can expect to be near the full allocation-value when there are at least several agents participating), the *redistribution mechanism* that I proposed in [2] requires, in aggregate, close to zero payments from the agents.

In the years since writing this first redistribution paper I have found that its core insight has application to a host of connected problems, which have come to constitute significant research projects in their own right. I briefly describe some of these here.<sup>1</sup>

*Efficient and budget-balanced allocation:* In single-item allocation the redistribution mechanism requires only small payments, but ideally it would require *no* payments. This is impossible to achieve with a mechanism that always allocates the good to the agent with highest value, but a variant of the redistribution mechanism requires no aggregate payments and only *very rarely* leads to misallocation of the good [6]. In trade settings, where a good is potentially to be *re-allocated*,

---

<sup>1</sup>[2] also spurred a mini-research area in the broader community, with papers on redistribution appearing frequently in the major AI and electronic commerce conferences, and occasionally in major economics journals, since its publication.

efficient allocation is impossible in equilibrium without running a deficit. I showed that a variant of the redistribution mechanism comes *closest* to efficiency [1].

*Auctions with entry and positive externalities:* The typical auction model starts with a pre-defined set of participants, but a richer approach models participation endogenously. Specifically, participation may be *costly*, and rational agents will participate only if their expected utility from doing so is non-negative. In a scenario where consumers are less inclined to bid in a low-participation (“obscure”) auction—imagine an upstart competitor to eBay trying to overcome its reputation and familiarity gap—the second-price auction is inefficient. In [4] I characterize efficient auctions for such settings, and show that a variant of the redistribution mechanism always improves efficiency. In other words, returning portions of the auction revenue back to the participants can increase the overall welfare by eliciting greater participation.

*Fairness:* In many cases, for a proposed decision procedure to be adopted by those with a stake in the outcome, it must be viewed as fair. In division of goods, this may entail restrictions on the extent to which any given agent would prefer the allocation obtained by others (envy), or restrictions to schemes in which each agent obtains at least a certain fraction of the utility that could be obtained as a dictator (proportionality). There is a large body of work on the fair division of goods, with many possibility and impossibility results regarding such fairness criteria. In contrast, mechanism design with monetary transfers is often oriented towards aggregate social welfare while ignoring other considerations. In [8] I join these distinct fairness and welfare goals: we would like to reach decisions that are deemed both fair and socially valuable. Interestingly, in addition to its superior welfare performance, it turns out that the redistribution mechanism is also remarkably *fair* (yielding *low* envy and *high* proportionality).

## Dynamic mechanism design

In my dissertation I extended the mechanism design enterprise to *dynamic* settings, those in which a sequence of decisions is to be made over time. For example, imagine a scenario 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. For instance in online display advertising, an advertiser’s value for advertising to a given web user may depend on whether they have done so already in the past.

There are two dimension to this kind of problem: a *planning* dimension of the sort studied extensively in the artificial intelligence community, which can be characterized as a Markov decision process (MDP); and an *incentives* problem related to the self-interest of agents combined with the private nature of their reward information, which is critical to forming the appropriate MDP representation. 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 real settings, though, agents often *will* acquire new private information over time. In [13], in collaboration with Prof. David Parkes (Harvard University) and also Prof. Satinder Singh (University of Michigan) 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. In [14] we applied these ideas to a semi-autonomous robot coordination problem.

I subsequently expanded and generalized this theory in several ways. In [3], I exactly characterized the entire class of dynamic mechanisms that are social-welfare maximizing in equilibrium, and in [15] I proposed a mechanism that is robust to a changing population of agents or scenarios in

which agents are periodically “inaccessible” to the mechanism. The proposed mechanism can be applied, for instance, to cases in which communication between agents and the social planner is faulty (communication breakdowns lead to intermittent inaccessibility). Interestingly, this work also unifies the complementary theories of online and dynamic mechanism design under one framework. Further work I’ve done in dynamic mechanism design includes:

*Dynamic redistribution:* In [3] I addressed the design of an efficient mechanism that requires only small payments in a *dynamic* decision setting. The dynamic case is significantly trickier than the static case discussed in the previous section because payments made in any given period, if dependent on information previously reported by an agent, can produce unintended incentives for manipulation. I specified a dynamic redistribution mechanism applicable to any setting that can be modeled as a multi-armed bandits problem, including, for instance, settings of repeated allocation of a single item. As in the static case, the proposed mechanism maintains the vast majority of value within the group of agents (i.e., payments to the central planner are very small).

*Metadeliberation auctions:* In [12], in collaboration with David Parkes I applied some of these recent theoretical advances in dynamic mechanism design to the problem of coordinating execution of an optimal policy in a resource allocation setting where agents can potentially “improve” their value for the resource, at a cost and with uncertain results, by performing some research or deliberation action. The planner must trade-off the possible gains of agent deliberation with the costs of waiting to make an allocation decision and the cost of the deliberation itself, essentially solving a multi-armed bandit planning problem.

## The human dimension

The point of mechanism design is to engineer good outcomes when *human* interests clash. With the rise of the internet it is now not uncommon for aspects of these interactions to take place by proxy via software agents; while this may expand the scope and (possibly) the complexity of the mechanisms that can be deployed, it does not diminish the centrality of the human interests that the agents exist to serve. Thus an appropriate model of human behavior is critical to any potential solution. Because behavior is so complex and varied, some simplifying assumptions and generalizations will be required in order to make progress, but this is a sensitive business because the resulting theory will be fragile to the extent that the adopted models are not predictive.

*Group decision-making when beliefs are subjective:* One salient case of “assumption overreach” has to do with belief formation in the presence of uncertainty. The classic economic model is that of perfect Bayesian updating, with an added assumption that all individuals start with *common prior* beliefs. An implication of this is that once groups of agents share all privately held information they will have identical posterior expectations—they cannot publicly “agree to disagree”—a prediction that often clashes with observed realities. In [5], drawing from the experimental psychology literature, I instead adopt a *subjective* belief model under which agents may indeed continue to disagree after sharing private information. This gives rise to a unique mechanism design problem, to which I provide a solution.

*Altruism:* Another ubiquitous assumption in mechanism design is that agents are indifferent to the utilities obtained by others. If instead agents are even *mildly* altruistic—discounting but not completely disregarding others’ interests—then the classic “solutions” break down. In [7] I introduce a regret-based model of altruism and examine its implications for mechanism design. To summarize one of the results: small amounts of altruism can have great effect in facilitating social-welfare maximizing decisions regarding resource allocation.

## Internet application areas

During my time at industry-based research organizations (Yahoo and Microsoft) I’ve worked on a number of practical problems that are highly amenable to the tools I’ve been describing.

*Display advertising auctions:* One prominent example is the case of internet display advertising, wherein web publishers sell space on their pages to advertisers who bid to display an ad image. This domain gives rise to a unique set of challenges. To name one, inventory is often sold “per-impression” (i.e, per page *view*), although there are some advertisers who only obtain value—and are only willing to pay—if their advertisement is *clicked*; buying an impression is risky for them because they do not know how likely it is that the viewer will click on the ad. Meanwhile there are arbitrage agents that make it their business to predict whether a given viewer will click on a given ad, and who are willing to serve as intermediaries, buying per-impression and then selling (per-click) to the advertiser. Their participation adds value to the system by facilitating efficient allocation, and introduces an important mechanism design question: can we derive a payment scheme that maximizes the expected value of the inventory, which involves getting all parties to reveal their private information, and yet does not diminish the revenue to the publisher compared to what would be obtained in its absence? I addressed this question in collaboration with Preston McAfee and Sergei Vassilvitskii [11], and validated our approach on a large dataset of bids from Yahoo’s Right Media display advertising exchange.

*Crowdsourcing procurement:* The internet has made connecting producers and consumers possible with an unprecedented level of facility and granularity. “Crowdsourcing” refers to the phenomenon of interaction between individuals and a large, typically anonymous workforce. Of particular economic interest is crowdsourcing *procurement*, wherein a principal solicits a certain type of good, and then multiple producers independently and simultaneously present submissions in pursuit of a prize offered to the producer of the good the principal deems of highest quality. In collaboration with Shaili Jain, I was the first to address the design of a social welfare-maximizing marketplace for this process [9]. In follow-up work [10], we study the efficiency level of the current prize schemes typical of the most prominent online crowdsourcing marketplaces.

## II Some future directions

The impetus for much of my research to date has been to make mechanism design more effective and applicable in the real world, even though many of the results are fundamental and theoretical. In the future I want to continue this “pragmatic theory” approach in new ways, which I believe is particularly critical for a field that aspires to help organize *human* interactions. While I am pursuing next steps to most of the projects described in the previous section, I am also charting new directions.

### Data-based design

Particularly for interactions that take place online, as the scope of data collection and its availability mushrooms, there is less and less need for guesswork in modeling agent behavior. For instance when a display advertising marketplace considers rolling out a new auction mechanism, it can be iteratively pilot-tested on a small fraction of the inventory—which may constitute *millions* of data points—prior to a full-scale roll-out. In other domains, or when the goal is not as quantifiable as pure revenue-maximization, things are less simple, but this new data-rich world offers a huge opportunity for mechanism design that is driven by machine learning of *revealed* behavioral realities rather than idealized models.

## Computational limitations: a chink in the armor

One of the most important assumptions lurking in standard mechanism design theory is that a socially optimal decision can be tractably computed; this assumption is unreasonable in most realistic dynamic settings and also in important static settings. This is a serious problem that has the potential to undermine the conclusions of a great portion of mechanism design theory, and it's a place where the computer science community holds the potential to make a large contribution. We need a framework with which to formalize a set of weaker "sufficient conditions" for cooperation in environments where computing optimal policies and strategies—or even precisely evaluating any proposed heuristic—is intractable. Behavioral analysis will play a large role here.

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<sup>2</sup> provides an efficient solution to the multi-armed bandit learning/decision-making problem; this result, in turn, implies the feasibility of an optimal dynamic mechanism for repeated allocation of a single item. The import of Gittins's achievement is hard to overstate, yet many significant closely related questions 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.

## Nearly efficient mechanisms

To go along with the glorious positive results mechanism design has mapped out over the past 30 years or so, there are some dispiritingly strong negative results: certain basic goals are in principle simply unachievable. Perhaps most notably, the Myerson-Satterthwaite theorem<sup>3</sup> states that for trade settings with selfish agents there can be no mechanism that is social-welfare maximizing and doesn't run a deficit in expectation. In the face of this negative result, a natural question is: how bad are its implications *in practice*? In a working paper [1], I start to demonstrate that in resource allocation settings the situation may not be very bad at all. Specifically, I adapt the intuition of the redistribution mechanism [2] to design a novel mechanism that never runs a deficit, and "almost always" provides the proper incentives to reach efficiency. The mechanism is not an approximation optimized for the worst-case, which may be uninformative about performance in expectation; rather, given a set of natural constraints such as the no-deficit restriction, it is the "least suboptimal" trading mechanism in *every case*.

This work-in-progress, like the altruism work described in the previous section, is part of a broader agenda: to take many of the core results of mechanism design (particularly the negative ones) and see whether they are as impactful in the real world as they are in the Platonic world of strict assumptions and hyper-idealized success criteria. This primarily involves design work, demonstrating *constructively* that new approaches may succeed in places where the idealized negative results have perhaps stifled efforts toward innovation.

## Distributing computation and decision-making

Multi-agent decision settings often provide natural opportunities for distributed computation. For instance, in a multi-agent sequential planning problem with private information, in each period one may ask each agent to report its local MDP, and then centrally compose and solve the multi-agent

---

<sup>2</sup>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.

<sup>3</sup>R. Myerson and M. A. Satterthwaite. Efficient mechanisms for bilateral trading. *Journal of Economic Theory*, 28:265–281, 1983.

MDP to determine an optimal joint-policy (specifying each agent’s decision). But it may also be possible for agents to locally contribute to *computation* of the optimal policy.

When self-interest is involved things get tricky, as agents may have incentive to shirk computing responsibilities or misreport the results of computation, and also privacy issues must be considered. In some problem domains independence relations mean optimal-policy computation is inherently factored, making the prospect of distributed solutions more compelling. In my work on *metadeliberation auctions*, I identified a mechanism in which agents, rather than simply reporting private information, will honestly compute and report the *Gittins index* associated with their local MDPs, which then allows the central planner to identify the optimal policy simply by comparing the reported indices. However, for a host of natural (and practical) closely related variant problems, the methods employed in that work do not lead to computationally tractable distributed solutions. To name one, consider a setting where a resource of uncertain value is to be allocated and agents can *learn better estimates* of their values through a costly sampling process. How can a central planner coordinate the learning process in a way that optimally trades off the costs of this learning with the expected gains in allocation value that it brings?

Another area where distributing computation brings associated incentives issues is *online crowd-sourcing procurement*, discussed above. If the production process is extended in time with incremental signals of the quality being produced, under certain conditions—which depend on the results of *others’* production—an agent should cease producing (i.e., if it becomes sufficiently unlikely that her produced good will be of the highest quality). What needs to be communicated across the agents, and what incentives need to be provided, to realize a globally-efficient production profile?

These examples are part of an agenda to apply mechanism design to multi-agent systems and electronic commerce applications in a way that goes 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.

## References

- [1] Ruggiero Cavallo. Optimal participation incentives for efficient and budget-balanced trade. Presented at the Decentralization Conference, Pasadena, CA, 2012.
- [2] Ruggiero Cavallo. Optimal decision-making with minimal waste: Strategyproof redistribution of VCG payments. In *Proceedings of the 5th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-06)*, pages 882–889, 2006.
- [3] Ruggiero Cavallo. Efficiency and redistribution in dynamic mechanism design. In *Proceedings of the 9th ACM Conference on Electronic Commerce (EC-08)*, pages 220–229, 2008.
- [4] Ruggiero Cavallo. Improving allocations through revenue redistribution in auctions with entry. In *The Second Conference on Auctions, Market Mechanisms and Their Applications (AMMA-11)*, New York, NY, 2011.
- [5] Ruggiero Cavallo. Incentives in group decision-making with uncertainty and subjective beliefs. In *Conference on Uncertainty and Artificial Intelligence (UAI-11)*, Barcelona, Spain, 2011.
- [6] Ruggiero Cavallo. Strongly budget-balanced and nearly efficient allocation of a single good. In *Workshop on Bayesian Mechanism Design, EC-11*, San Jose, CA, 2011.
- [7] Ruggiero Cavallo. Strongly budget-balanced auctions with mildly altruistic agents. In submission, 2011.

- [8] Ruggiero Cavallo. Fairness and welfare through redistribution when utility is transferable. In *Proceedings of the 26th Annual Conference on Artificial Intelligence (AAAI-12), Toronto, Canada, 2012*.
- [9] Ruggiero Cavallo and Shaili Jain. Efficient crowdsourcing contests. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS-12), Valencia, Spain, 2012*.
- [10] Ruggiero Cavallo and Shaili Jain. Winner-take-all crowdsourcing contests with stochastic production. In *Workshop on Social Computing and User Generated Content, EC-12, Valencia, Spain, 2012*.
- [11] Ruggiero Cavallo, Preston McAfee, and Sergei Vassilvitskii. Display advertising auctions with arbitrage. In *Workshop on Ad Auctions, EC-12, Valencia, Spain, 2012*.
- [12] Ruggiero Cavallo and David C. Parkes. Efficient metadeliberation auctions. In *Proceedings of the 26th Annual Conference on Artificial Intelligence (AAAI-08)*, pages 50–56, 2008.
- [13] 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.
- [14] Ruggiero Cavallo, David C. Parkes, and Satinder Singh. Optimal coordination of loosely-coupled self-interested robots. In *Workshop on Auction Mechanisms for Robot Coordination, AAAI-06, Boston, MA, 2006*.
- [15] Ruggiero Cavallo, David C. Parkes, and Satinder Singh. Efficient mechanisms with dynamic populations and dynamic types. Under review, early version presented at the DIMACS Workshop on the Boundary Between Economic Theory and Computer Science, June 2011.