Economic reasoning and artificial intelligence

Citation

Parkes, D. C., and M. P. Wellman. 2015. "Economic Reasoning and Artificial Intelligence." *Science* 349 (6245) (July 16): 267–272. doi:10.1126/science.aaa8403. <http://dx.doi.org/10.1126/science.aaa8403>.

Published Version

doi:10.1126/science.aaa8403

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:25622971>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Economic Reasoning and Artificial Intelligence

David C. Parkes¹
Michael P. Wellman²

¹School of Engineering and Applied Sciences, Harvard University,
33 Oxford St., Cambridge MA 02138, USA

²Computer Science and Engineering, University of Michigan,
2260 Hayward St, Ann Arbor, MI 48109, USA

The field of artificial intelligence (AI) strives to build rational agents, capable of perceiving the world around them and taking actions to advance specified goals. Put another way, AI researchers aim to construct a synthetic *homo economicus*, the mythical perfectly rational agent of neoclassical economics. We review progress towards creating this new species of machine, *machina economicus*, and discuss some challenges in designing AIs that can reason effectively in economic contexts. Supposing that AI succeeds in this quest, or at least comes close enough that it is useful to think about AIs in rationalistic terms, we ask how to design the rules of interaction in multi-agent systems that come to represent an economy of AIs. Theories of normative design from economics may prove more relevant for artificial agents than human agents, with AIs that better respect idealized assumptions of rationality than people, interacting through novel rules and incentive systems quite distinct from those tailored for people.

Rationality in Economics and in AI

Economics models the behavior of people, firms, and other decision makers, as a means to understand how these decisions shape the pattern of activities that produce value and ultimately satisfy (or fail to satisfy) human needs and desires. In this enterprise, the field classically starts from an assumption that actors behave rationally, that is, their decisions are the best possible given their available actions, their preferences, and their beliefs about the outcomes of these actions. Economics is drawn to rational decision models because they directly connect choices and values, in a mathematically precise manner. Critics argue that the field studies a mythical species, *homo economicus* (“economic man”), and produces theories with limited applicability to how real humans behave. Defenders acknowledge that rationality is an idealization, but counter that the abstraction supports powerful analysis, which is often quite predictive of people’s behavior (as individuals, or in aggregate). Even if not perfectly accurate representations, rational models also allow preferences to be estimated from observed actions and build understanding that can usefully inform policy.

Artificial Intelligence research is likewise drawn to rationality concepts, as they provide an ideal for the computational artifacts it seeks to create. Core to the modern conception of AI is the idea of designing agents: entities that perceive the world and act in it (*I*). The quality of an AI design is judged by how well the agent’s actions advance specified goals, conditioned on the perceptions observed. This coherence among perceptions, actions, and goals is the essence of rationality. If we represent goals in terms of preference over outcomes, and conceive perception and action within the framework of decision making under uncertainty, then the AI agent’s situation aligns squarely with the standard economic paradigm of rational choice. Thus, the AI designer’s task is to build rational agents, or agents that best approximate rationality given the limits of their computational resources (2–4). In other words, AI strives to construct—out of

silicon (or whatever) and information—a synthetic *homo economicus*, perhaps more accurately termed *machina economicus*.

The shared rationality abstraction provides a strong foundation for research that spans AI and economics. We start this review by describing progress on the question of how to operationalize rationality, and how to construct AI agents that are able to reason about other AIs. Supposing that AI research succeeds in developing an agent that can be usefully modeled as rational (perhaps more so than human agents), we turn to research on the design of systems populated by multiple AIs. These multi-agent systems will function as AI economies, with AIs engaged in transactions with other AIs as well as with firms and people. This prospect has spawned interest in expanding theories of normative design from economics, optimizing rules of encounter (5) to guide multi-agent interactions. Systems populated by AIs may exhibit new economic phenomena, and thus require a new science with which to understand the way they function and to guide their design. For example, while human cognitive constraints limit the design of current markets, systems designed for AIs may admit more complex interfaces, impose greater calculation burdens, and demand more stamina of attention.

At the same time, the ways in which the behavior of AIs deviate from the behavior of people can present new challenges. We can already glimpse the future of economic AIs, with simple AI bots pricing books for sale on Amazon and scanning for restaurant tables on Opentable for re-sale at a profit (6). Such AIs may introduce some efficiencies, but their lack of common sense and their designer’s failure to anticipate interactions can also lead to books priced at \$23m (7). More sophisticated AI strategies, presumably more carefully vetted, exert a large influence on financial markets, with automated trading algorithms estimated to be responsible for more than 70% of trades on U.S. stock markets (8). Given the consequences, it is important to understand the impact of ubiquitous automated agents on the performance of economic systems. As reasoning is shifted from people to AIs—designed to learn our preferences, overcome our decision

biases, and make complex cost-benefit tradeoffs—how too should the economic institutions that mediate everyday transactions change?

We focus in this review on some of the research directions we consider most salient for a future synthesis of economics and AI engendered by the emergence of *machina economicus*. Interesting as they are, we only briefly mention here the many exciting applications of AI to problems in economics such as matching (9), market clearing (10) and preference modeling for smart grids (11). Nor will we showcase the many ways in which economic theory is finding application today within AI, for example game-theoretic approaches to multi-agent learning (12) and voting procedures to combine the opinions of AIs (13).

Building *Machina Economicus*

Constructing a rational AI raises a host of technical challenges not previously addressed in the long tradition of rationalistic modeling in the social sciences. For economics, the agent attitudes (e.g., beliefs, preferences) underlying rationality are conceptual abstractions. Economists need not explain how capabilities and preferences, for example, are encoded, nor the algorithm by which an agent plans what actions to take conditional on its perceptions. Computation is abstracted away in the standard economic model, and is precisely what the AI scientist must account for in order to operationalize rationality in a realized agent.

This does not mean an AI design needs to incorporate data structures corresponding directly to rationality constructs, though many AI architectures do feature direct representations for propositions, goals, and the like. Such representations may simplify the analysis of AI systems, for example we can ask whether an inference algorithm operating on logical expressions possesses desirable properties such as soundness: that all conclusions follow from the premises. Similarly, if an AI's beliefs are encoded as probability distributions, we can ask whether it updates its beliefs from observations in proper accord with Bayesian theory. However, care must

Figure 1: A bounded reinforcement learning agent performs better by pursuing a designed reward function different from the *objective reward*: its actual fitness evaluation. Results (right) from a gridworld foraging domain (left), for various limits on the agent’s planning horizon (16). Unless the agent is perfectly rational (i.e., no horizon limit)—not typically feasible in realistic applications—the designer can often achieve better fitness by directing the agent to optimize an alternative measure.

be taken in understanding an agent’s attitudes solely in terms of its internal data structures. Imperfections in decision making may mean that the beliefs held and objectives pursued by a computational agent, in effect, vary systematically from those directly encoded.

As an example illustrating this distinction, machine learning researchers adapted from animal learning the concept of reward shaping (14). In reinforcement learning, the agent derives a policy (mapping from perception sequences to actions) based on rewards representing instantaneous value associated with a state and action. A designer specifying the input reward can often train the agent more efficiently by shaping the reward signal over the learning process, to facilitate convergence to behavior optimizing the designer’s objective. The framework of optimal rewards (15) provides a general treatment distinguishing reward specifications and designer goals. As shown in Fig. 1, the optimal reward input to the agent does not generally correspond to the designer’s ideal reward. This perspective helps explain the role of intrinsic motivations (e.g., curiosity) in a flexible learning agent.

Though the mantle of designing *machina economicus* may not be adopted (particularly in such explicit terms) by all AI researchers, many AI advances over the last few decades can be characterized as progress in operationalizing rationality. For instance, probabilistic reasoning was largely eschewed by AI thirty years ago, but now pervades the field, thanks to developments in representation and inference using Bayesian networks and related graphical formalisms. Expressing uncertainty about general relationships, beyond mere propositions, is routinely supported in probabilistic modeling languages (17). Statistical approaches now dominate machine learning and natural language processing [cite review article in this issue]. Likewise preference handling (including methods for eliciting preferences from the designer of an AI agent, compactly representing preferences over complex domains, and enabling inference about preferences) is regarded as a necessary AI facility. Planning, the AI subfield concerned with action over time, now conventionally frames its problem as one of optimization, subject to resource constraints, multiple objectives, and probabilistic effects of actions.

Will AI succeed in developing the ideal rational agent? As much as we strive to create *machina economicus*, absolutely perfect rationality is unachievable with finite computational resources. A more salient question is whether AI agents will be sufficiently close to the ideal as to merit thinking about them and interacting with them in rationalistic terms. Such is already the case, at least in a limited sense. Whenever we anthropomorphize our machines, we are essentially treating them as rational beings, responding to them in terms of our models of their knowledge, goals, and intentions. A more refined version of the question is whether our formal rationality theories will fit well the behavior of AI agents, in absolute terms, or compared to how well the theories work for people. Without offering any judgment on the question of how well rationality theories capture essential human behavior, we note the irony in the prospect that social science theories may turn out to apply with greater fidelity to non-human agent behavior.

Reasoning about other Agents

The issue of agent theorizing is not merely academic. If we can build one AI agent, then we will build many, and these AIs will need to reason about each other, as well as about people. For AIs designed to approximate *machina economicus*, it stands to reason they should treat each other as rational, at least as a baseline assumption. These AIs would adopt a game-theoretic view of the world, where agents rationally respond to each others' behavior, presumed (recursively) to be rational as well. A consequence is that agents would expect their joint decisions to be in some form of equilibrium, as in standard economic thinking.

That AIs (or AI-human combinations) are reasonably modeled as approximately rational is the premise of a growing body of AI research applying economic equilibrium models to scenarios involving multiple agents (18). The approach has achieved notable successes, providing evidence for the premise, at least in particular circumstances. Just as single-agent rationality does not require literal expected-utility calculations, applicability of an equilibrium model does not require that agents themselves be explicitly engaged in equilibrium reasoning. For example, the literature on learning in games (19) has identified numerous conditions where simple adaptive strategies converge to strategic equilibria. We can evaluate the effectiveness of economic modeling by examining agents built by AI designers for specified tasks. For instance, in a study of AI trading agents competing in a shopping game (20), an agent using standard price equilibrium models from economics (specifically, Walrasian equilibrium) achieved comparable prediction accuracy to sophisticated machine learning approaches, without using any data—even though none of the other agents employed equilibrium reasoning.

In the rest of this section, we describe further examples where economic modeling, in the form of game-theoretic algorithms, has provided an effective way for AIs to reason about other agents.

The first example is computer poker. Although poker is an artificial game, many humans have invested a great deal of time and money to develop their playing skills. More importantly, poker's uncertainty and complexity have made it a compelling challenge problem for AI techniques. Early approaches aimed to capture the knowledge of expert human players (21), but over the last decade, game-theoretic algorithms have predominated. Technically, poker is a game of imperfect information, where each player knows elements of history (cards dealt to them) that are secret from others. As uncertainty gets partially resolved over time, through card turns and betting, players must update their beliefs, about both card outcomes and the beliefs of others.

A major milestone in computer poker was achieved in 2014, with the effective solution of heads up limit hold'em (HULHE), which is a standard two-player version of the most popular poker game (22). HULHE is the largest game of imperfect information ever solved (over 10^{13} information sets, after removing symmetries), and the first imperfect-information game widely played by humans to be solved. The solution was the culmination of two decades of effort by a series of researchers (See Fig. 2), beginning with exact solution of simplified poker games, and proceeding to approximate solution of abstracted versions of the full game (23). Computing the approximate Nash equilibrium of the full game required massive computation and new methods for equilibrium search based on regret-matching techniques from machine learning. The result is a strategy against which even a perfect opponent cannot earn a detectable profit.

In general, the optimal strategy against perfect opponents may not be the ideal strategy against the more typical fallible kind. Despite considerable effort, however, researchers have not found poker algorithms that perform considerably better than game-theoretic solutions, even against natural distributions of opponents. It has also turned out that game-theoretic approaches have been more successful than alternatives, even for poker variants that are far from being exactly solved, such as no-limit (where bets are unrestricted) (24), or games with three or more

Figure 2: Researchers produced steady exponential progress on solving games of imperfect information from 1995 to the present. Up to 2007 (left), game size was generally reported in terms of nodes in the game tree. Based on methods introduced around that time, it became more meaningful (right) to report size in terms of the number of information sets (each many nodes), which represent distinct situations as perceived from the perspective of a player. The circled data points correspond to the same milestone; combining the two graphs thus demonstrates the continual exponential improvement.

players (25).

Much of the interest in game-theoretic reasoning for AI is driven by its applicability to real-world problems. The most prominent area of application in recent years and our second example is that of security games, based on a pioneering series of systems developed by Tambe et al. (26). In these systems, an agent decides how to defend facilities (e.g., airport security through placement of checkpoints) by solving a game where an attacker is presumed to rationally plan in response to the defender’s decision. This approach has been successfully deployed in a variety of domains, including airport and airline security and coast guard patrols.

As for any game-theoretic approach, the recommendations from these systems are sensitive to assumptions made about the other agents (here, attackers): their respective preferences, beliefs, capabilities, and level of rationality. Representational approaches from AI provide flexibility, allowing the assumptions made in the strict versions typically employed by game theorists

to be relaxed (27). The field of behavioral game theory has developed detailed predictive models based on how humans have been observed to deviate from game-theoretic rationality (28). Such predictive models can be readily incorporated in existing game-theoretic reasoning algorithms, as has been demonstrated in the context of modeling attackers in security games (29). An interesting open question is whether the kinds of behavioral models that best explain human decision making (see Wright and Leyton-Brown (30) for a meta-study) will also prove effective in capturing the bounded rationality of computational agents.

Designing Multi-Agent Systems

At the multi-agent level, a designer cannot directly program behavior of the AIs, but rather defines the rules and incentives that govern interactions among AIs. The idea is to change the “rules of the game” (e.g., rewards associated with actions and outcomes) in order to effect change in agent behavior and achieve system-wide goals. System goals might include, for instance, promoting an allocation of resources to maximize total value, coordinating behavior in order to complete a project on time, or pooling decentralized information to form an accurate prediction about a future event. The power to change the interaction environment is special, and distinguishes this level of design from the standard AI design problem of performing well in the world as given.

An interesting middle ground is to take the world as given, but employ reliable entities—mediators—that can interact with AIs and perform actions on their behalf (31). Introducing mediating entities is relatively straightforward in the new AI economy. To see how this can be powerful, consider a mediated extension of the classic Prisoner’s dilemma game (Fig. 3). If both AIs grant the mediator the authority to play on their behalf (i.e., proxy right), it performs Cooperate on behalf of both agents. However, if only one AI grants the mediator proxy, it performs Defect on behalf of that agent. In equilibrium, both AIs grant proxy, and the effect

	Cooperate	Defect		Mediator	Cooperate	Defect
Cooperate	4,4	0,6		Mediator	4,4	6,0
Defect	6,0	1,1		Cooperate	0,6	4,4
				Defect	1,1	6,0
			(a)			(b)

Figure 3: Each entry gives the utility to (row player, column player). (a) Prisoner's dilemma. The dominant strategy equilibrium is (Defect, Defect). (b) Mediated Prisoner's dilemma. The dominant strategy equilibrium is (Mediator, Mediator).

is to change the outcome from (Defect,Defect) to (Cooperate,Cooperate), increasing utility to both participants.

For the more general specification of rules of interaction for rational agents, economics has a well-developed mathematical theory of mechanism design (32). The framework of mechanism design has been fruitfully applied for example to the design of matching markets (33) and auctions (34). Mechanism design is a kind of inverse game theory, with the rules inducing a game, and the quality of the system evaluated in an equilibrium. In the standard model, design goals are specified in terms of agent preferences on outcomes, but these preferences are private and agents self-interested. A mechanism is a trusted entity, able to receive messages from agents that make claims (perhaps untruthfully) about preferences, and select an outcome (e.g., an allocation of resources, a plan of behavior) on the basis of these messages. The challenge is to align incentives and promote truthful reports.

Varian (35) has argued that the theory of mechanism design may actually prove more relevant for artificial agents than human agents, as AIs may better respect the idealized assumptions of rationality made in this framework. For example, one desirable property of a mechanism is incentive compatibility, which stipulates that truthful reports constitute an equilibrium. Sometimes it is even possible to make truthful reporting a dominant strategy (optimal whatever others do), achieving the strong property of strategy-proofness (36). But it seems that people do not reliably understand this property; evidence from medical matching markets, and also from lab

experiments, suggests that some participants in strategy-proof matching mechanisms try to misrepresent their preferences even though it provides no advantage (37, 38).

For artificial systems, in comparison, we might expect AIs to be truthful where this is optimal, and to avoid spending computation reasoning about the behavior of others where this is not useful (5). More generally, mechanism designs for AI systems need not be simple because they need not be understandable to people. On the contrary, AI techniques such as preference representation, preference elicitation, and search algorithms can be used to turn the mathematical formalisms of mechanism design into concrete computational methods (39–41). The design problem itself can also be usefully formulated as a computational problem, with optimization and machine learning used to find solutions to design problems for which analytical solutions are unavailable (42–45).

The prospect of an economy of AIs has also inspired expansions to new mechanism design settings. Researchers have developed incentive-compatible multi-period mechanisms, considering such factors as uncertainty about the future and changes to agent preferences because of changes in local context (46–48). Another direction considers new kinds of private inputs beyond preference information (49, 50). For example, in a team formation setting, each AI might misreport information about the capabilities of other AIs in order to get itself selected for the team (51). Similarly, AIs seeking to maximize task assignments might provide false reports of experience in task performance in order to mislead a learning mechanism constructing an automatic task classifier (52). Systems of AIs can also create new challenges for mechanism design. One such challenge is false-name bidding, where an AI exploits its ability to manage multiple identities. For example, it may gain resources more cheaply by dividing a request into a set of smaller requests, each placed from a different identity under its control. In response, researchers have developed mechanisms that are robust to this new kind of attack (53).

The important role of mechanism design in an economy of AIs can be observed in practice.

Search engines run auctions to allocate ads to positions alongside search queries. Advertisers bid for their ads to appear in response to specific queries (e.g., “personal injury lawyer”). Ads are ranked according to bid amount (as well as other factors such as ad quality), with higher ranked ads receiving a higher position on the search results page. Early auction mechanisms employed first-price rules, charging an advertiser its bid amount when its ad receives a click. Recognizing this, advertisers employed AIs to monitor queries of interest, ordered to bid as little as possible to hold onto the current position. This practice led to cascades of responses in the form of bidding wars, amounting to a waste of computation and market inefficiency (54). To combat this, search engines introduced second-price auction mechanisms (36), which charge advertisers based on the next highest bid price rather than their own price. This approach (a standard idea of mechanism design) removed the need to continually monitor the bidding to get the best price for position, thereby ending bidding wars (55).

In recent years, search engine auctions have supported richer, goal-based bidding languages. For example, an advertiser can ask to maximize clicks over a weighted set of queries subject to a budget constraint (56, 57). Search engines provide proxy agents that then bid on behalf of advertisers to achieve the stated goal (58). This introduction of proxy agents and the earlier switch from first price to second price can be interpreted as a computational application of a fundamental concept in mechanism design—the revelation principle (59–61). Briefly, this states that if the rules of a mechanism and the equilibrium strategies in that mechanism are replaced by a new mechanism that is functionally equivalent to the composition of these rules and strategies, then the new mechanism will be incentive compatible. Although neither redesign provides incentive compatibility in a formal sense, both second-pricing and proxy bidding can be interpreted as accomplishing on behalf of advertisers what they were doing (through AIs) in an earlier design. See Fig. 4. Still other ad platform designs are using a strategy-proof mechanism (the Vickrey-Clarke-Groves mechanism (36, 62, 63)) to make decisions about the

Figure 4: Two generations of sponsored search mechanisms. Early designs were first price (FP) and advertisers (ADV) used AIs (AI-POS) to maintain a position on the list of search results at the lowest possible price. Second price (SP) auction mechanisms were introduced, designed to replace the combination of FP and AI-POS. Advertisers adopted new AIs (AI-GOAL) to achieve higher-level goals such as maximize profit or maximize the number of clicks. The second price auction was extended to include proxy agents (SP+Proxy), designed to replace the combination of SP and AI-GOAL.

space to allocate to ads, which ads to allocate, and which (non-sponsored) content to display to a user (64).

The tangle between automated agents and the design of rules of interaction also features prominently in today's financial markets, where the dominance of computerized traders has, by most accounts, qualitatively shaped the behavior of these markets. Although details of implementation are closely held secrets, it is well understood that techniques from AI and machine learning are widely employed in the design and analysis of algorithmic traders (65). Algorithmic trading has enabled the deployment of strategies that exploit speed advantages, and has led in turn to a costly arms race of measures to respond to market information with minimum latency. A proposed design response would replace continuous-time auctions with periodic auctions that clear on the order of once per second, thus negating the advantage of tiny speed improvements (66, 67).

We describe two additional examples of the design of multi-agent systems for an economy of AIs.

The first example system aggregates information held by multiple AIs. The rules of a system

that achieves this goal can be engineered purposefully through the design of a prediction market (68). Popular versions of prediction markets feature questions such as who will be elected U.S. President (e.g., Betfair offers many such markets). The basic idea of a prediction market is to facilitate trade in securities contracts (e.g., a possible contract will pay \$1 if Hilary Clinton elected). The price that balances supply and demand is then interpreted as a market prediction (e.g., price \$0.60 reflects probability 0.6 for the payoff event).

Consider a domain with a large number of interrelated random variables, for example “flight BA214 delayed by more than 1 hour”, “snowstorm in Boston”, “de-icing machine fail”, “incoming flight BA215 delayed by more than 1 hour”, and “security alert in London”. In a combinatorial prediction market (69), a large bet on the contract “de-icing machine fail” would affect the price of “flight BA214 delayed by more than 1 hour”, and all other connected events. A challenge is that the number of conceivable events is exponential in the number of random variables. Among other properties, a good market design should allow bets on all events about which AIs have information (e.g., “de-icing machine fail AND all subsequent flights from Boston delayed by more than 1 hour”.) A good design should also align incentives, for example making it utility-maximizing to trade immediately on current information until the market price reflects an agents belief. Progress in scaling up combinatorial markets has been made by relating the problem of pricing bets to well-understood problems in statistical inference and convex optimization (70, 71). Related research advances are being made by allowing AIs to transact in hypotheses that are acquired through machine learning as well as trade directly in information signals rather than beliefs (72–74).

The second example is the management of information concerning the trustworthiness of agents within an economy of AIs. Trust that a counterparty will complete a transaction or invest effort or resources is crucial for any well-functioning economic system. A standard approach is to associate participants with a reputation, which can serve to align incentives in the present

under the threat of a damaged reputation and lost opportunities in the future. In addition to this problem of moral hazard (i.e., will agents behave cooperatively when completing economic transactions), reputation systems can address the problem of adverse selection (i.e., will high quality agents choose to enter a market in the first place) (75, 76).

A special challenge in an economy of AIs arises because of the fluidity of identity, and the ease with which agents can be replaced. This raises, for example, the spectre of whitewashing attacks, where an AI repeatedly runs down its reputation before re-entering anew with a different identity. Without the possibility of enforcing strong identities that cannot be changed, this suggests a social cost of fluid identities, where it becomes necessary to impose a penalty on all new participants and make them build up reputation from an assumption of being untrustworthy (77).

We should also consider that *machina economicus* will be strategic in sharing feedback on other AIs. For example, in eBay's original reputation system, buyers were often reluctant to leave negative feedback about deadbeat sellers, because the sellers could retaliate with negative feedback about the buyer. In response, eBay introduced an additional feedback mechanism that was one-directional from the buyer to the seller and could not be easily traced to a particular buyer. The change resulted in a greater amount of negative feedback (78).

The economy of AIs also offers positive opportunities for promoting trust through book-keeping, collecting feedback, and for tracking the provenance of feedback in novel reputation mechanisms (see Fig. 5). AI researchers are designing reputation systems that align incentives with making truthful reports, while provably satisfying axiomatic properties such as symmetry: two agents that are in an equivalent position from the perspective of reports made and received should have the same trust score (79, 80). Another example is the design of accounting systems that elicit truthful reports about the resources contributed or work performed by other AIs, and enable the design of systems to mitigate free-riding and promote fair contributions to

Figure 5: In a reputation system for a multi-agent AI, each agent chooses an action, and the combined effect of these actions generates rewards (i.e., utility). Based on the actions taken and rewards received, agent i can submit a report, x_i to the reputation system. The reputation system aggregates this feedback, for example providing a ranked list to reflect the estimated trustworthiness of agents. Each agent observes this ranked list, and this information may influence future actions.

an economic system (81). Still, the extent to which effective, multi-agent AIs can be developed entirely through computational infrastructure such as reputation mechanisms and without recourse to legal systems remains an interesting open question.

Closing Comments

Whatever one's thoughts about when or whether AI will transcend human-level performance, the rapidly advancing capabilities of AI are fueling significant optimism and investment in AI research. AI has surpassed or will likely soon surpass humans in narrow domains such as playing chess, controlling a jumbo jet during cruise, making product recommendations, pricing

millions of products on an eCommerce platform, reasoning about whether a patient is likely to be re-admitted to a hospital, and detecting signals from a massive volume of financial news stories.

Certainly, many fundamental challenges remain, including how to design reasoning and inference methods that effectively balance the benefit of additional computation with the costs that may arise from additional delay to acting in the world, and how to design AI systems that can learn and generalize from reward signals in unconstrained domains. But given that decision problems related to economic transactions are often relatively well structured, it seems likely to us that AI will continue to make especially rapid inroads in economically significant applications. This in turn will ensure continued effort on methods for rational, economic reasoning, toward the broader goal of developing *machina economicus*.

We should not leave the impression that AI researchers unanimously embrace economic perspectives on single- or multi-agent AI. For some, multi-agent economic models are still seen as a distraction. After all, a centralized perspective allows focusing on overall goals without worrying about the incentives of individual parts of the system. Others conduct research into multi-agent systems comprised of agents under the control of the designer, so that they can be programmed in any way desired. Just as with centralized solutions, these so-called “co-operative” multi-agent systems allow design without concern for the self-interest of individual agents, albeit often with decomposition or communication constraints. But cooperative versus self-interested is really a difference in assumptions on the power of a system designer, rather than a technical dispute. The viewpoint that we ascribe to is that a large number of AI systems will, given the existing structure of human economic systems, be populated by AIs that are designed, deployed, owned and operated by a myriad of different parties each with possibly misaligned goals. Finally, some may object to the economic approach on the basis that AIs are and will remain far from perfectly rational, simply by virtue of physical and computational

limits. More direct models of the AIs' computational behavior, in terms of the automata they are, could in principle be more accurate. The analytical utility of a rationality abstraction for AIs is ultimately an empirical question to be resolved as AI progresses.

Among those adopting an economic approach, there persist some disagreements on specific techniques, for example on the role of equilibrium reasoning. Even if agents can be viewed as rational, some question whether it is plausible that they reach equilibrium configurations, particularly in situations where multiple equilibria exist. As Shoham (82) argues, game theory lacks a well-accepted pragmatic account of how it should be deployed in concrete reasoning contexts. A positive view is that AI researchers, in their efforts to operationalize economic reasoning, are developing exactly this needed body of pragmatics.

Some may object that mechanism design is too idealized even for systems of AIs, for example in its insistence on design under equilibrium behavior, its assumption that rules of interaction can be designed from scratch, and its lack of attention to the details of the human and legal contexts in which designed systems will operate. A positive view is that AI systems are precisely the kinds of environments where we can build *tabula rasa* new rules of interaction since these rules will be realized through the Internet and as programs running on computer servers. That rules of interaction can come into existence is as much a matter of science and engineering as it is of public policy.

As AI advances, we are confident that economic reasoning will continue to have an important role in the design of single-agent and multi-agent AIs, and we have argued that as economies of AIs continue to emerge there will need to be a new science to understand how to design these systems. These AIs will no doubt exert strong forces on the economy and broader society; understanding the effect and extent of this will shape the research agendas of both AI and economics in years to come.

References

1. S. Russell, P. Norvig, *Artificial Intelligence: A Modern Approach* (Prentice Hall, 2009), third edn.
2. J. Doyle, *Computational Intelligence* **8**, 376 (1992).
3. E. J. Horvitz, Computation and action under bounded resources, Ph.D. thesis, Stanford University (1990).
4. S. Russell, *Fundamental Issues of Artificial Intelligence*, V. C. Müller, ed. (Springer, 2015).
5. J. S. Rosenschein, G. Zlotkin, *Rules of Encounter: Designing Conventions for Automated Negotiation among Computers* (MIT Press, 1994).
6. M. McCue, Can't secure a reservation to a hot restaurant? Try using an algorithm (2014). Downloaded from <http://fortune.com/2014/09/03/opentable-dinematic-tablesweep-restaurant-reservation-service> on June 16, 2015.
7. J. D. Sutter, Amazon seller lists book at \$23,698,655.93—plus shipping (2011). Downloaded from <http://edition.cnn.com/2011/TECH/web/04/25/amazon.price.algorithm> on June 16, 2015.
8. T. Hendershott, C. M. Jones, A. J. Menkveld, *Journal of Finance* **66** (2011).
9. D. J. Abraham, A. Blum, T. Sandholm, *Eighth ACM Conference on Electronic Commerce* (2007), pp. 295–304.
10. A. Fréchette, N. Newman, K. Leyton-Brown, *24th International Joint Conference on Artificial Intelligence* (2015).

11. S. D. Ramchurn, P. Vytelingum, A. Rogers, N. R. Jennings, *Communications of the ACM* **55**, 86 (2012).
12. R. I. Brafman, M. Tennenholtz, *Artificial Intelligence* **159**, 27 (2004).
13. A. X. Jiang, *et al.*, *Advances in Neural Information Processing Systems* (2014), pp. 2573–2581.
14. A. Y. Ng, D. Harada, S. J. Russell, *16th International Conference on Machine Learning* (1999), pp. 278–287.
15. S. Singh, R. L. Lewis, A. G. Barto, J. Sorg, *IEEE Transactions on Autonomous Mental Development* **2**, 70 (2010).
16. J. Sorg, S. Singh, R. Lewis, *27th International Conference on Machine Learning* (2010), pp. 1007–1014.
17. L. Getoor, B. Taskar, eds., *Introduction to Statistical Relational Learning* (MIT Press, 2007).
18. Y. Shoham, K. Leyton-Brown, *Multiagent Systems: Algorithmic, Game-Theoretic and Logical Foundations* (Cambridge University Press, 2009).
19. D. Fudenberg, D. K. Levine, *The Theory of Learning in Games* (MIT Press, 1998).
20. M. P. Wellman, D. M. Reeves, K. M. Lochner, Y. Vorobeychik, *Journal of Artificial Intelligence Research* **21**, 19 (2004).
21. D. Billings, A. Davidson, J. Schaeffer, D. Szafron, *Artificial Intelligence* **134**, 201 (2002).
22. M. Bowling, N. Burch, M. Johanson, O. Tammelin, *Science* **347**, 145 (2015).

23. T. Sandholm, *AI Magazine* **31**, 13 (2010).
24. N. Brown, S. Ganzfried, T. Sandholm, *14th International Conference on Autonomous Agents and Multi-Agent Systems* (2015), pp. 7–15.
25. R. Gibson, Regret minimization in games and the development of champion multiplayer computer poker-playing agents, Ph.D. thesis, University of Alberta (2014).
26. M. Tambe, *Security and Game Theory: Algorithms, Deployed Systems, Lessons Learned* (Cambridge University Press, 2011).
27. Y. Gal, A. Pfeffer, *Journal of Artificial Intelligence Research* **33**, 109 (2008).
28. C. F. Camerer, *Behavioral Game Theory: Experiments in Strategic Interaction* (Princeton University Press, 2003).
29. R. Yang, C. Kiekintveld, F. Ordóñez, M. Tambe, R. John, *Artificial Intelligence* **195**, 440 (2013).
30. J. R. Wright, K. Leyton-Brown, *CoRR* **abs/1306.0918** (2013).
31. D. Monderer, M. Tennenholtz, *Artificial Intelligence* **173**, 180 (2009).
32. L. Hurwicz, *Mathematical Methods in the Social Sciences*, K. J. Arrow, S. Karlin, P. Suppes, eds. (Stanford University Press, 1960).
33. T. Sönmez, U. Ünver, *Handbook of Social Economics*, A. Bisin, J. Benhabib, M. Jackson, eds. (North-Holland, 2011), vol. 1A, pp. 781–852.
34. P. Milgrom, *Putting Auction Theory to Work* (Cambridge University Press, 2004).
35. H. R. Varian, *First USENIX Workshop on Electronic Commerce* (1995), pp. 13–21.

36. W. Vickrey, *Journal of Finance* **16**, 8 (1961).
37. Y. Chen, T. Sönmez, *Journal of Economic Theory* **127**, 202 (2006).
38. F. Echenique, A. J. Wilson, L. Yariv, Clearinghouses for two-sided matching: An experimental study, *Working Papers 487*, University of Pittsburgh, Department of Economics (2013).
39. N. Nisan, *Combinatorial Auctions*, P. Cramton, Y. Shoham, R. Steinberg, eds. (MIT Press, 2006), chap. 9.
40. T. Sandholm, C. Boutilier, *Combinatorial Auctions*, P. Cramton, Y. Shoham, R. Steinberg, eds. (MIT Press, 2006), chap. 10.
41. T. Sandholm, *Artificial Intelligence* **135**, 1 (2002).
42. V. Conitzer, T. Sandholm, *18th Conference on Uncertainty in Artificial Intelligence* (2002), pp. 103–110.
43. S. Alaei, H. Fu, N. Haghpanah, J. D. Hartline, A. Malekian, *13th ACM Conference on Electronic Commerce* (2012), p. 17.
44. Y. Cai, C. Daskalakis, S. M. Weinberg, *54th Annual IEEE Symposium on Foundations of Computer Science* (2013), pp. 618–627.
45. P. Duetting, *et al.*, *ACM Transactions on Economics and Computation* **3** (2015).
46. D. C. Parkes, S. Singh, *Advances in Neural Information Processing Systems* (2003).
47. R. Cavallo, D. C. Parkes, S. Singh, *22nd Conference on Uncertainty in Artificial Intelligence* (Cambridge, MA, 2006), pp. 55–62.

48. D. C. Parkes, *Algorithmic Game Theory*, N. Nisan, T. Roughgarden, E. Tardos, V. Vazirani, eds. (Cambridge University Press, 2007), chap. 16, pp. 411–439.
49. Y. Shoham, M. Tennenholtz, *Theoretical Computer Science* **343**, 97 (2005).
50. J. Y. Halpern, V. Teague, *36th Annual ACM Symposium on Theory of Computing* (2004), pp. 623–632.
51. N. Alon, F. Fischer, A. Procaccia, M. Tennenholtz, *13th Conference on Theoretical Aspects of Rationality and Knowledge* (2011), pp. 101–110.
52. O. Dekel, F. A. Fischer, A. D. Procaccia, *Journal of Computer and System Sciences* **76**, 759 (2010).
53. M. Yokoo, Y. Sakurai, S. Matsubara, *Games and Economic Behavior* **46**, 174 (2004).
54. B. Edelman, M. Ostrovsky, *Decision Support Systems* **43**, 192 (2007).
55. B. Edelman, M. Ostrovsky, M. Schwarz, *American Economic Review* **97**, 242 (2007).
56. C. Borgs, *et al.*, *World-Wide Web Conference* (2007), pp. 531–540.
57. J. Feldman, S. Muthukrishnan, M. Pál, C. Stein, *Eighth ACM Conference on Electronic Commerce* (2007), pp. 40–49.
58. A. Z. Broder, E. Gabrilovich, V. Josifovski, G. Mavromatis, A. J. Smola, *Proceedings of the Forth International Conference on Web Search and Web Data Mining, WSDM 2011, Hong Kong, China, February 9-12, 2011* (2011), pp. 515–524.
59. L. Hurwicz, *Decision and Organization*, R. Radner, McGuire, eds. (North-Holland, Amsterdam, 1972).

60. A. Gibbard, *Econometrica* **41**, 587 (1973).
61. R. Myerson, *Econometrica* **47**, 61 (1979).
62. E. Clarke, *Public Choice* **11**, 17 (1971).
63. T. Groves, *Econometrica* **41**, 617 (1973).
64. H. R. Varian, C. Harris, *American Economic Review* **104**, 442 (2014).
65. M. Kearns, Y. Nevmyvaka, *High Frequency Trading: New Realities for Traders, Markets and Regulators* (Risk Books, 2013).
66. E. Budish, P. Cramton, J. Shim, The high-frequency trading arms race: Frequent batch auctions as a market design response, *Tech. Rep. 14-03*, Booth School of Business, University of Chicago (2015).
67. E. Wah, M. P. Wellman, *14th ACM Conference on Electronic Commerce* (2013), pp. 855–872.
68. R. Forsythe, F. Nelson, G. R. Neumann, J. Wright, *Contemporary Laboratory Experiments in Political Economy*, T. R. Palfrey, ed. (University of Michigan Press, 1991), pp. 69–111.
69. R. D. Hanson, *Information Systems Frontiers* **5**, 107 (2003).
70. J. Abernethy, Y. Chen, J. W. Vaughan, *ACM Transactions on Economics and Computation* **1**, 12:1 (2013).
71. M. Dudík, S. Lahaie, D. M. Pennock, D. Rothschild, *14th ACM Conference on Electronic Commerce* (2013), pp. 341–358.
72. J. Abernethy, R. Frongillo, *Advances in Neural Information Processing Systems* (2011).

73. J. Witkowski, D. C. Parkes, *13th ACM Conference on Electronic Commerce* (2012), pp. 964–981.
74. J. Hu, A. J. Storkey, *31st International Conference on Machine Learning* (2014), pp. 1773–1781.
75. P. Resnick, K. Kuwabara, R. Zeckhauser, E. J. Friedman, *Communications of the ACM* **43**, 45 (2000).
76. C. Dellarocas, *Management Science* **49**, 1407 (2003).
77. E. J. Friedman, P. Resnick, *Journal of Economics & Management Strategy* **10**, 173 (2001).
78. G. Bolton, B. Greiner, A. Ockenfels, *Management Science* **59**, 265 (2013).
79. A. Cheng, E. Friedman, *ACM SIGCOMM Workshop on Economics of Peer-to-Peer Systems* (2005), pp. 128–132.
80. A. Altman, M. Tennenholtz, *Journal of Artificial Intelligence Research* **31**, 473 (2008).
81. S. Seuken, J. Tang, D. C. Parkes, *24th AAAI Conference on Artificial Intelligence* (2010), pp. 860–866.
82. Y. Shoham, *Communications of the ACM* **51**, 75 (2008).
83. Thanks to the editor, the anonymous referees, Michael Bowling, Yiling Chen, Kobi Gal, Sébastien Lahaie, David Pennock, Ariel Procaccia, Tuomas Sandholm, Sven Seuken, Yoav Shoham, Amos Storkey, Milind Tambe, and Moshe Tennenholtz for their thoughtful comments on an earlier draft.