

Designing Simple Mechanisms

Shengwu Li

Economists today design rules for real-world marketplaces. For example, economists in the 1990s designed auctions for radio spectrum broadcast licenses and reformed the National Resident Matching Program (NRMP), a clearinghouse that matches graduates of American medical schools to their first jobs (Milgrom 2000; Roth 2002). The practice soon found myriad applications; for instance, economists have designed systems to coordinate electricity generation, to match students to public schools, to swap donated goods between food banks, and to sell loans from the Bank of England to financial institutions (Wilson 2002; Abdulkadiroğlu et al. 2005; Prendergast 2017, 2022; Klemperer 2010).

In these mechanisms, participants convey information about their preferences, which affects who gets what and (sometimes) how much they pay. Participants will naturally try to figure out whether it makes sense to provide information in accordance with their true preferences, or whether it might prove advantageous to strategize—that is, whether the mechanism is “incentive-compatible.” If a mechanism is simple to play, then it will be easier for participants to see whether it is in their interest to reveal their true preferences.

But what makes a mechanism simple? When is it easy for participants to see that a mechanism is incentive-compatible? I will start here by explaining how and why economists came to ask these questions. Then I will discuss three recent answers, which will capture different aspects of what makes a mechanism simple.

■ *Shengwu Li is Professor of Economics, Harvard University, Cambridge, Massachusetts. His email address is shengwu_li@fas.harvard.edu.*

For supplementary materials such as appendices, datasets, and author disclosure statements, see the article page at <https://doi.org/10.1257/jep.38.4.175>.

Taking Mechanisms Literally

Early work in mechanism design treated mechanisms as metaphors. For instance, Hurwicz (1973) describes a “resource allocation mechanism” as follows:

Simplifying to the utmost, we may imagine each agent having in front of him a console with one or more dials to set; the selection of dial settings by all agents determines uniquely the flow of goods and services (trade vector) between every pair of agents and also each agent’s production (input-output vector), his “trade with nature.”

Hurwicz is describing an abstract representation of richer real-world institutions. The dials stand for purchases that the agents could make, bids that they could place, bargaining strategies that they could adopt, and so on. Later work converged on a standard abstraction, in the form of the “revelation principle”: Under some conditions, it is without loss of generality to restrict attention to truthful behavior in incentive-compatible revelation mechanisms. In such mechanisms, each agent is asked to report their private information directly, and it is in their best interest to report truthfully (Gibbard 1973; Dasgupta, Hammond, and Maskin 1979; Myerson 1979). A mechanism is *dominant-strategy incentive-compatible*, or “strategy-proof,” if reporting truthfully is always a best response, regardless of the behavior of other players.

On the metaphorical interpretation of a mechanism, it does not matter whether participants understand that the mechanism is incentive-compatible. After all, people do not interact with the mechanism *per se*; instead, the mechanism represents aspects of their everyday economic transactions. People may be proficient at arranging those transactions—making purchases in shops, haggling over contracts—and yet not recognize that the mechanism is incentive-compatible when confronted with its abstract form.

The practice of designing real-world markets led economists to focus on a more literal interpretation of mechanisms, namely that mechanism design involves “explicitly analyzing the consequences of trading rules that presumably are really common knowledge” (Wilson 1987). On this view, the rules of the mechanism are not shorthand for a vague decentralized process; rather, they capture actual rules of the institution as understood by participants. For instance, the rules for a sealed-bid auction specify what bids are permitted and how the winner and the payments are determined.

When we regard the mechanism as capturing the rules of a formal process, such as an auction or a school choice system, it matters whether participants understand that the mechanism is incentive-compatible. This is a pressing question for market design, for several reasons. First, market designers often consider novel design proposals that are unfamiliar to the intended participants—or worse, may be deceptively familiar, leading participants to adopt heuristics that are unsuited to the new rules. Second, in applications such as school choice or

certain high-stakes auctions, some people participate in the mechanism exactly once, and thus cannot learn from experience. Third, even if the designer asserts that the mechanism is incentive-compatible, participants may distrust that claim. Moreover, that distrust may be justified: before the intervention of economists, the National Resident Matching Program (mistakenly) claimed that its algorithm made it incentive-compatible for medical students to report their true preferences, even though this was false (Williams 1995).¹ Similarly, when Google started selling Internet advertising via auction, it incorrectly claimed that its auction format was strategy-proof, asserting that its “unique auction model uses Nobel Prize-winning economic theory to eliminate [. . .] that feeling you’ve paid too much” (Edelman, Ostrovsky, and Schwarz 2007). (This presumably referred to the 1996 prize to William Vickrey and James Mirrlees, for work that we will shortly discuss.) To the best of my knowledge, both claims were genuine mistakes rather than intentional deception.

Participants in real-world mechanisms do not always recognize that the mechanism is incentive-compatible. The modern algorithm for the National Resident Matching Program is (for all practical purposes) incentive-compatible,² and the program prominently advises that applicants should rank jobs in their preferred order. Nonetheless, 17 percent of applicants claimed, when surveyed, to have submitted a nontruthful rank-order list (Rees-Jones 2018). When Rees-Jones and Skowronek (2018) recruited medical students who recently participated in the National Resident Matching Program to take part in an incentivized experiment with the same algorithm, they found that 23 percent of subjects submitted nontruthful rank-order lists in the experiment. As further evidence, in strategy-proof mechanisms that match students to degree programs, some students make unambiguous mistakes. They rank a program without a scholarship above the same program with a scholarship, even though the scholarship is worth thousands of euros. This behavior has been found in matching mechanisms for graduate psychology degrees in Israel (Hassidim, Romm, and Shorrer 2021) and for undergraduate degrees in Hungary (Shorrer and Sívágó 2023).

Consequently, it is often not enough that a mechanism is theoretically incentive-compatible. Participants have to see for themselves that it is incentive-compatible; the mechanism has to be simple. Simplicity bypasses the need for participants to trust the mechanism designer. Simplicity eases the cognitive cost of participation, which is just as real as other costs. Simplicity can level the playing field, in the sense of Pathak and Sönmez (2008), by preventing unfair outcomes caused by unequal

¹At the time, the National Resident Matching Program ran the hospital-proposing deferred acceptance algorithm, which is strategy-proof for the hospitals listing their preferred candidates, but not incentive-compatible for the applicants themselves.

²The program now uses the applicant-proposing deferred acceptance algorithm, modified to account for couples. The standard deferred acceptance algorithm is strategy-proof for the proposing side (Dubins and Freedman 1981), but accounting for couples slightly undermines this property. Computational experiments indicate that the probability that an applicant can profitably deviate under the modern algorithm is about 1 in 10,000 (Roth and Peranson 1999).

strategic sophistication. Finally, simplicity allows the designer to rely more confidently on the predictions of classical game theory, because that analysis depends on participants responding predictably and correctly to incentives.³

In the rest of this article, I explain three challenges that participants may need to overcome, in order to recognize incentive-compatibility: (1) thinking contingently about unobserved moves by other players; (2) planning for their own future moves; and (3) reasoning about other players' beliefs. In each of these cases, I will explain formal criteria that capture the difficulty in question, and discuss mechanisms that alleviate that difficulty.

These three challenges are not exhaustive of the ways that mechanisms can be simple or complex. This article is not a comprehensive survey, and it is shaped by my idiosyncratic tastes and the limits of my expertise. I have omitted excellent papers for the sake of brevity.

Thinking Contingently about Unobserved Moves by Other Players

Some mechanisms require participants to reason case-by-case about other players' moves, in order to see that the mechanism is incentive-compatible, while others do not.

In a seminal paper, Vickrey (1961) invented the *second-price sealed-bid auction*,⁴ in which all participants simultaneously place bids, then the highest bidder wins and pays the second-highest bid. Vickrey argued that the second-price auction is desirable because it is simple: "Each bidder can confine his efforts and attention to an appraisal of the value the article would have in his own hands, at a considerable saving in mental strain and possibly in out-of-pocket expense."

The second-price sealed-bid auction is strategy-proof, so it is in your own best interest to bid your true value. One can see this by analogy. In a dynamic *ascending auction*, the price starts low and gradually rises. At each moment, each bidder decides whether to keep bidding or to quit irrevocably. When just one bidder is left, that bidder clinches the object at the current price. It is obvious that you should keep bidding if the price is below your true value, and should quit if the price is above. Vickrey pointed out that the second-price auction is "logically isomorphic to" the ascending auction. In the ascending auction, each bidder essentially chooses when to quit. The bidder with the highest quit-price wins, and pays the second-highest quit-price. Quit-prices in the ascending auction are equivalent to bids in the second-price auction, so it is a dominant strategy to place a bid equal to your value.

³Li and Dworzak (2024) make a formal critique of this last justification. If participants are strategically unsophisticated, then instead of using a simple mechanism, the designer might profitably adopt complex mechanisms that confuse the participants.

⁴It is more accurate to say that Vickrey reinvented the second-price auction. Second-price auctions were used to sell postage stamps to collectors as early as 1893 (Lucking-Reiley 2000).

Notice that Vickrey’s argument that the second-price sealed-bid auction is incentive-compatible works by first reasoning about a dynamic mechanism—the ascending auction—and then arguing that two mechanisms are equivalent. But do real participants treat these two auctions as equivalent?

For auctions in the wild, it can be hard to tell whether participants are bidding truthfully, because we usually do not know their value for the object. To overcome this obstacle, Kagel, Harstad, and Levin (1987) had lab subjects bid in auctions for a virtual prize, which had a different value for each bidder. All subjects knew their own value, in dollars, for the virtual prize. The winning bidder received cash equal to their value minus the price they paid. This method enables the experimenter to observe whether subjects are bidding truthfully. Kagell, Harstad, and Levin (1987) randomized whether subjects participated in second-price sealed-bid auctions or in dynamic ascending auctions. Subjects played their assigned auction format up to 30 times in a row, with their values drawn randomly each time. Again, second-price auctions and ascending auctions are equivalent in theory; it is a dominant strategy to choose a bid (or quit-price) equal to your value. Thus, we might expect that in both formats subjects will bid truthfully, and that the clearing price will always be equal to the second-highest value.

The theoretical prediction of truthful bidding was largely borne out in ascending auctions. Subjects rapidly figured out that truthful bidding is optimal, and in 76 percent of auctions the clearing price was at the second-highest value. By contrast, in second-price sealed-bid auctions, subjects did not behave as theory predicted. Subjects often submitted bids far from their values, and the clearing price was at the second-highest value in no more than 20 percent of auctions. Even after 30 rounds of play, prices remained far from the theoretical prediction.

In summary, Kagell, Harstad, and Levin (1987) found that lab subjects rapidly converge on truthful bidding in ascending auctions, but make large and persistent mistakes in second-price auctions. This result is well-replicated (for examples, see Kagel and Levin 1993; McCabe, Rassenti, and Smith 1990; Harstad 2000; Li 2017; Breitmoser and Schweighofer-Kodritsch 2022). Ascending auctions and second-price auctions are not equivalent in practice.

What did our theory miss? Why is it easy for people to see that the ascending auction is strategy-proof? Why is it hard to see this for the second-price auction?

Imagine a participant who is bidding in a second-price sealed-bid auction. Let us assume that this participant is not a trained economist, and thus does not realize that the second-price auction is “logically isomorphic to” the ascending auction.⁵ Suppose that the participant values the object at \$40. To see that bidding \$40 is better than deviating to a different bid, the participant must make a case-by-case comparison, calculating payoffs for each profile of opponent bids. In particular, the participant must understand that if a truthful bid of \$40 would win at a price

⁵In a lab experiment, Breitmoser and Schweighofer-Kodritsch (2022) found that subjects are more likely to bid truthfully in second-price auctions when the mechanism is framed in a way that makes the isomorphism salient. But why does the ascending auction framing make it easier to see the dominant strategy?

of \$30 (because the highest other bid is \$30), then deviating to bid \$50 would also lead to a win at \$30. And if a truthful bid of \$40 would lose, then a bid of \$50 could only win at prices above \$40. If the participant does not keep track of each contingency, then they might be tempted to bid above their true value in order to raise the chance of winning. Of course, this strategy would be a mistake, because boosting their bid to \$50 raises their chance of winning only when the highest other bid is between \$40 and \$50. In that case, the price is so high that winning is undesirable.

Now imagine a participant who is bidding in a dynamic ascending auction, and values the object at \$40. The price is now \$30. If the participant continues to bid truthfully, then the participant will either win at a price between \$30 and \$40, or will lose and have payoff zero. Thus, every possible payoff from bidding truthfully is at least zero. On the other hand, if the participant deviates to quit right now, then their payoff is zero for sure.⁶ The worst-case payoff from truthful bidding is at least as good as the best-case payoff from deviating, so the participant need not keep track of different contingencies to see that truthful bidding is optimal.

A symmetric argument establishes that the participant should quit when the price exceeds their value. Suppose that the participant values the object at \$40 and the price rises to \$41. If the participant quits, then their payoff is zero for sure. If instead the participant keeps bidding, then their payoff can never be positive, because now they can only win at prices above \$41.

In summary, we can distinguish dynamic ascending auctions and second-price sealed-bid auctions in this way: To play optimally in a second-price auction, you need to compare strategies case-by-case, whereas to play optimally in an ascending auction, you need not make that comparison.

In Li (2017), I generalized this idea to any extensive-form mechanism. A strategy is *obviously dominant* if, for any deviating strategy, at any information set where the two strategies choose different actions for the first time, the worst-case payoff under the dominant strategy is at least as good as the best-case payoff under the deviation. (When the worst case from the dominant strategy is preferred to the best case from the deviation, a case-by-case comparison is not needed.) A mechanism is *obviously strategy-proof* if the truthful strategy is obviously dominant.

Dynamic ascending auctions are obviously strategy-proof. In contrast, at the moment you submit your bid in a second-price sealed-bid auction, bidding your value can lead to a zero payoff (if you lose), whereas bidding above your value can yield a positive payoff (depending on the second-price bid). Thus, truthful bidding is not an obviously dominant strategy, and second-price auctions are not obviously strategy-proof.

Every obviously strategy-proof mechanism is strategy-proof. In fact, a stronger claim holds: Every obviously strategy-proof mechanism is *weakly group strategy-proof*, meaning that no coalition of players can jointly deviate and all strictly benefit. We will

⁶We have ignored the possibility of ties. Formally, we are modeling the ascending auction as a mechanism in which we cycle between the active bidders in some fixed order, asking each to raise their bid by a dollar or to quit irrevocably. When just one bidder remains, that bidder wins and pays their last bid.

prove the contrapositive. Suppose the mechanism is not weakly group strategy-proof. Then there exists a strategy profile and a deviating coalition, such that the deviation is strictly profitable for every member. Along the resulting path-of-play, consider the first coalition member to deviate from the truthful strategy. That player strictly benefits, so at that information set, one possible outcome from the deviating strategy is strictly better than one possible outcome from the truthful strategy. It follows that the truthful strategy is not obviously dominant, which completes the proof.

Obvious strategy-proofness characterizes the ascending auction. Suppose that we restrict attention to auctions that are efficient (allocating the object to a bidder with maximal value) and in which only the winning bidder makes payments. In that class, an extensive-form mechanism is obviously strategy-proof if and only if it is an ascending auction (Li and Lo 2023, pp.576–82). Under weaker restrictions, one can characterize obviously strategy-proof mechanisms as “personal-clock auctions,” a larger class that includes ascending auctions, descending-price reverse auctions, and some hybrid formats (Li 2017). Roughly, in a personal-clock auction, each participant chooses between clinching a sure thing (for example, leaving empty-handed) and a tentative alternative with a price that gets steadily worse for the participant (for example, bidding to buy the object at a price that increases each round). For another example, consider a reverse auction in which the bidder is a seller, facing a gradually descending price. At each point, the seller can either quit and keep their object (the sure thing) or offer to sell at the current price (the tentative alternative).

In a practical application of these issues, the US Federal Communications Commission (FCC) auctioned off \$19.8 billion of radio spectrum rights in 2017. It included a reverse auction that paid television stations to relinquish their over-the-air broadcast rights, so that this spectrum could be repurposed for wireless Internet access. The Federal Communications Commission’s reverse auction had to deal with complicated feasibility constraints. To check whether an allocation is feasible, one must solve a graph coloring problem for a graph with up to 3,000 nodes and 2.7 million edges; the nodes represent television stations and the edges represent interference constraints. Stations that share an edge cannot be assigned to the same channel. This problem is computationally hard (Karp 1972). Consequentially, the reverse auction used a suite of state-of-the-art feasibility-checking algorithms to determine the allocation and payments (Leyton-Brown, Milgrom, and Segal 2017).

The reverse auction proceeded, roughly, as follows: The Federal Communications Commission makes an opening offer to each television station. Each television station can either quit or continue. If it quits, then it remains on the air, retaining its broadcast rights with zero net payment. If it continues, then either it sells its rights at the current price, or the FCC makes a lower offer and the process repeats. It is an obviously dominant strategy for each television station to bid truthfully, quitting once the price drops below its value for remaining on the air. This reverse auction belongs to the class of deferred-acceptance clock auctions studied by Milgrom and Segal (2020), which have useful implications for privacy and budget-balance.

I have not described how the reverse auction determines the offers or when it decides to close a sale. Those details depend on the complicated feasibility-checking algorithms, and most bidders lacked the computational horsepower to check whether the algorithms were running as intended. But even without those details, you can see that the reverse auction is incentive-compatible. This is no accident. A mechanism is obviously strategy-proof if and only if truth-telling can be seen to be dominant from a certain kind of partial description. Such a description specifies the sequences of queries that one participant might receive, the answers that are allowed, and how queries and answers map to possible outcomes (Li 2017, Theorem 1). Notice that this description omits how the outcome depends on unobserved opponent moves; in this sense, obviously strategy-proof mechanisms formalize the idea of a dominant strategy that can be recognized without contingent reasoning.

Psychologically, it is harder to account for hypothetical contingencies than to reason about observed events (Esponda and Vespa 2014). Dynamic mechanisms can help participants to avoid mistakes, by resolving uncertainty about what other players have done.⁷

The assumptions we make about preferences affect the structure of obviously strategy-proof mechanisms. Recent work has explored beyond auctions, characterizing obviously strategy-proof mechanisms for settings such as two-sided matching, social choice, and object allocation without transfers (Arribillaga, Massó, and Neme 2020, 2023; Ashlagi and Gonczarowski 2018; Bade 2019; Bade and Gonczarowski 2017; Mandal and Roy 2022; Thomas 2020; Troyan 2019). Other work has built tools to study obviously strategy-proof mechanisms, providing a “revelation principle” (Mackenzie 2020) and algorithms to construct them (Golowich and Li 2021).

Planning for Future Moves

Some mechanisms require participants to plan far in advance for their own future moves, in order to see that the mechanism is incentive-compatible, while other mechanisms do not require such planning.

Suppose that we have a set of indivisible goods and a set of agents, and we want to allocate one good to each agent, without money transfers. One natural solution is to order the agents randomly, give the first agent their favorite good, then give the second agent their favorite good among those that remain, and so on iteratively. This algorithm results in efficient allocation: any other allocation that makes one agent better-off must also make another agent worse-off (Bogomolnaia and Moulin 2001).

⁷This insight connects to a larger literature on contingent reasoning (Martínez-Marquina, Niederle, and Vespa 2019; Cohen and Li 2022; Esponda and Vespa 2023), reviewed by Niederle and Vespa (2023). Glazer and Rubinstein (1996) formalize another sense in which dynamic games can be easier to understand than their static counterparts.

To carry out the above algorithm, we need to ask agents about their preferences. We could approach agents one-by-one in random order, asking each to pick their favorite good among those that remain. The resulting *dynamic random priority mechanism* is obviously strategy-proof. In contrast, we could ask each agent to submit rank-order lists of the goods, and then process those lists according to the algorithm. The resulting *static random priority mechanism* is strategy-proof. However, it is not obviously strategy-proof. Because of the randomness of the order in which the rank-order lists are chosen, it is possible that if you rank goods truthfully, you might receive your third-favorite, whereas if you deviate to rank your second-favorite good at the top, then you might receive your second-favorite good, which you strictly prefer.

In practice, dynamic random priority results in higher rates of truthful play than its static equivalent. In an incentivized laboratory experiment with four players and four goods, dynamic random priority resulted in the dominant-strategy outcome in 93 percent of games, whereas static random priority resulted in the dominant-strategy outcome in only 64 percent of games (Li 2017, p. 3282).⁸

Ascending auctions and dynamic random priority are both intuitively simple. But some games with obviously dominant strategies are not intuitively simple. For example, take any chess position such that White can force a win, and consider the subgame that starts from that position. If White plays the win-forcing strategy, then by definition the worst-case outcome is that White wins. If White ever deviates from that strategy, then the best-case outcome is that White wins. Thus, the win-forcing strategy is obviously dominant in the subgame. However, real chess players often fail to play the win-forcing strategy, even from positions where such a strategy has been found by computers (Anderson, Kleinberg, and Mullainathan 2017). (One of the unanswered questions of chess is whether White can force a win from the opening position. Chess great Bobby Fischer once opined, “I think it’s almost definite that the game is a draw theoretically.”)

What makes dynamic random priority simple, and chess complicated? Intuitively, to identify a win-forcing strategy in chess, one has to plan far in advance, consider different future contingencies, and then backward-induct to judge the merits of each current move. By contrast, when playing dynamic random priority, each player chooses exactly once, so forward planning is not required. Conventional game theory elides this distinction, because a strategy is viewed as a complete contingent plan, specifying not only what one does right now, but also what one will do at all future contingencies.

One way that a mechanism can be simple is that it might require only limited forward planning. Pycia and Troyan (2023) formalized the idea of limited forward planning. They consider games of perfect recall, that is, the information sets are such that each participant “remembers” all the past information sets that they encountered and all the past moves they took. However, at each information set, the active player can only foresee some of their other information sets. Thus,

⁸Dreyfuss, Heffetz, and Rabin (2022) and Meisner and von Wangenheim (2023) argue that this disparity might be due to loss aversion rather than strategic mistakes.

each player forms a *partial strategic plan* based on their own information set, that specifies moves at foreseeable information sets. This plan is simply dominant if the worst-case outcome from following the plan is at least as good as the best-case from deviating to another plan that chooses differently at that information set. Specifically, the worst-case and the best-case are with respect to the actions of other players and also the actions that the player takes at unforeseeable information sets.

This approach to modeling limited forward planning results in a variety of incentive criteria, which depend on how we specify the foreseeable information sets. At one extreme, suppose that all information sets are foreseeable; in this case, simple dominance is equivalent to obvious dominance. At the other extreme, suppose that only the present information set is foreseeable, and let us call the resulting criterion *strong obvious dominance*. For example, in dynamic random priority, picking your favorite object is strongly obviously dominant. Each player is called to play exactly once, so it does not matter that future information sets are unforeseeable. Strong obvious dominance may seem too demanding a criterion, but there is a large literature studying the welfare and revenue guarantees of posted-price mechanisms, which have strongly obviously dominant strategies (Lucier 2017).

However, one can imagine an intermediate case, where the player can foresee some, but not all, of the future information sets. Consider a bidder who values the object at \$40. The bidder is participating in an ascending auction, with the price rising by \$1 in each step. The price starts at \$1, and the bidder reasons, “If I quit now my payoff is \$0 for sure. But if I agree to bid \$1 and plan to quit at \$2, then my payoff could be $\$40 - \$1 = \$39$, and is at least \$0. So I’ll keep going for now.” Then the price rises to \$2, and the bidder thinks, “If I quit now my payoff is \$0 for sure. But if I keep bidding, planning to quit at \$3, then my payoff could be $\$40 - \$2 = \$38$, and is at least \$0.” Thus, by looking just one step ahead at each point, the bidder is led to behave in a way that reproduces the optimal strategy. Notice that the bidder is making a collection of partial strategic plans, and these plans are not consistent with each other. When the price is \$1, they plan to quit at \$2. But when the price reaches \$2, the bidder revises that plan.

To formalize this idea, let the foreseeable information include both the current information set (which includes all the player’s earlier moves) and every information set that is “one step ahead.” Observe that the ascending auction is one-step simple, meaning that truthful bidding can be induced by a collection of partial strategic plans (one for each information set), each of which is simply dominant when the player looks just one step ahead. Just a little bit of foresight is enough to play optimally in an ascending auction.

Let us return to the allocation of indivisible goods without transfers. As we saw, dynamic random priority is obviously strategy-proof. But there are other obviously strategy-proof mechanisms in this setting; Pycia and Troyan (2023) characterized those mechanisms, and some of them are not intuitively simple, instead requiring participants to plan far into the future in order to see that the mechanism is incentive-compatible. They showed that strengthening the simplicity requirement

from obvious strategy-proofness to one-step simplicity rules out the counterintuitive mechanisms.

It can take detailed contingent thinking to see that a static mechanism, such as the second-price sealed-bid auction, is incentive-compatible. Dynamic mechanisms, such as the ascending auction, can mitigate this difficulty by paring down the contingencies that participants must consider. But dynamic mechanisms can raise new difficulties, in some cases requiring participants to plan in advance—even far in advance—to understand that a given strategy is optimal. The concept of partial strategic plans formally introduced by Pycia and Troyan (2023) provides a way to study this kind of complexity.⁹

Reasoning about Other Players' Beliefs

The study of simplicity is not only about refining strategy-proofness. Even mechanisms without dominant strategies can vary in how simple they are. In particular, some mechanisms require participants to reason about other players' beliefs, in order to see that the mechanism is incentive-compatible, while other mechanisms do not.

Consider bilateral trade with transfers; there is a seller who can produce an indivisible object, and a prospective buyer. The seller has a privately-known cost C , and the buyer has a privately-known value V . If a sale is made at some price p , then the buyer's payoff is $V - p$ and the seller's payoff is $p - C$. Otherwise, both payoffs are 0.

One natural mechanism for bilateral trade is the double auction, studied by Chatterjee and Samuelson (1983): The seller and buyer simultaneously submit offers, s and b respectively. They transact if and only if the offer s from the seller is less than or equal to the offer b from the buyer. The resulting price is a weighted average of s and b , that is, $\alpha s + (1 - \alpha)b$ for parameter $\alpha \in [0, 1]$. Suppose for the moment that $\alpha = 0.5$, so the mechanism splits the difference between the two offers. In this case, neither side has a dominant strategy. The seller's utility-maximizing offer depends on the buyer's offer b ; if the seller knew the buyer's offer b , then the seller would set $s = b$ if the seller's cost satisfies $C \leq b$, and set $s > b$ otherwise.

Suppose it is common knowledge that both players are rational, that is, they both make offers that maximize expected utility. Because both offers are made at the same time, the seller does not know the buyer's offer b . Instead, the seller has to form beliefs about the distribution of b . Because the seller knows that the buyer is rational, the seller knows that the buyer's choice of b depends on the buyer's value V and on the buyer's beliefs about the seller's offer s . And that, in turn, depends on the buyer's beliefs about the seller's cost C , and the buyer's beliefs about the seller's beliefs about b . Thus, the seller's offer depends on a

⁹Jéhiel (1995) and Jéhiel and Samet (2007) also study limited forward planning in games, but it is an open question whether one can adapt those ideas for use in mechanism design.

second-order belief; what the seller believes that the buyer believes about C , as well as a third-order belief; what the seller believes that the buyer believes that the seller believes about V , and so on.

This many-level reasoning process seems fantastical, and in practice will be beyond the capacity of many participants. If participants are inexperienced, it seems unlikely that they will find the equilibrium of the double auction with parameter $\alpha = 0.5$ from first principles. However, some mechanisms do not require participants to reason about higher-order beliefs. Suppose that we take the double auction, and instead set the split-the-difference parameter α equal to 1. Then the seller is effectively making a take-it-or-leave-it offer of s to the buyer, and the transaction occurs at a price of s if and only if $s \leq b$. In this case, it is a dominant strategy for the buyer to set $b = V$, so the buyer does not need to think about the seller's strategy. If the seller knows that the buyer will set $b = V$, then the seller's problem reduces to choosing an offer s to maximize $(s - C)$ multiplied by the probability that $s \leq V$. Observe that the solution depends on the seller's cost C and on the seller's belief about the buyer's value V , but not on any higher-order beliefs. The seller can calculate an optimal offer using only first-order beliefs, although the seller does not have a dominant strategy.

Börgers and Li (2019) proposed a new incentive criterion, picking out mechanisms in which participants can “play well” using just their first-order beliefs. The proposal is as follows: Let us fix, for each player, a set of possible utility functions (which can be thought of as the different values that each buyer might have for an object). Each player will then form a strategy, which means choosing an action based on their utility function; for example, in the double auction, the actions are the feasible offers. A strategy is *undominated* if there is no other strategy that is always at least as good, and sometimes strictly better. A *first-order belief* is a distribution over the utility functions of the other players; for instance, in a double auction, the seller might have a first-order belief that the buyer's value is uniformly distributed on all integers between 0 and 10. A strategy is robust with respect to that first-order belief if, assuming the other players' utilities are indeed distributed according to that belief, the strategy is a best response to every profile of undominated strategies for the other players. A mechanism is *strategically simple* if for every player and every first-order belief, there exists a robust strategy.

Every strategy-proof mechanism is strategically simple, but the reverse does not hold true. For example, the double auction with split-the-difference parameter $\alpha = 1$ is strategically simple, but not strategy-proof. For the buyer, offering $b = V$ is a dominant strategy, and hence a robust strategy. Moreover, offering $b = V$ is the buyer's only undominated strategy, so the seller has a robust strategy which depends only on the seller's cost C and on the seller's beliefs about the distribution of V (that is, the seller's first-order belief). By contrast, the double auction with $\alpha = 0.5$ is not strategically simple. The buyer has many undominated strategies, and for nontrivial seller beliefs there is no robust seller strategy.

Börgers and Li (2019) formally proved that in the bilateral trade setting, all the strategically simple mechanisms have a special structure: The “leading” player either chooses an offer p from a set of permitted prices, or declines trade. If an offer of p

is made, then the “following” player either accepts (and trade occurs at price p) or declines, in which case no trade occurs. This class includes the double auction with $\alpha = 0$ and with $\alpha = 1$, as well as variations that restrict the set of permitted prices.

More generally, Börgers and Li (2019) found that, under a richness condition, the strategically simple mechanisms can be characterized as “local dictatorships.” Roughly, this means that if we fix a profile of utility functions, and consider the restricted mechanism that includes only actions that are undominated *at that profile*, then in the restricted mechanism only one of the player’s actions determines the outcome. The meaning of what they call “local dictatorships” is subtle, and includes mechanisms that are not “dictatorial” in the colloquial sense. One example of a local dictatorship is a voting mechanism in which each committee member submits a list ranking all alternatives and (at the same time) the chairperson selects two alternatives for a head-to-head vote. From the selected pair of outcomes, the final outcome is the alternative that is ranked higher by a majority of the committee (suppose that all players have strict preferences over the alternatives). In this situation, each committee member has only a single undominated action, namely, to submit their truthful rank-order list. Thus, in the restricted mechanism, the chairperson’s action determines the outcome.

Conclusion

Mechanisms can be complicated in many ways, challenging players to think contingently about the moves of other players, to plan for their own future moves, and to reason about other players’ beliefs. Thus, there are multiple ways to define what makes a mechanism “simple.” The right criterion depends on context; after all, participants in designed mechanisms include schoolchildren, doctors, lumberjacks, fishermen, and telecommunications firms advised by game theorists (Athey and Levin 2001; Marszalec, Teytelboym, and Laksá 2020; Bulow, Levin, and Milgrom 2009).

Why have formal simplicity criteria at all? After all, we know simplicity when we see it. If we use intuition to judge formal criteria, why not use intuition instead of formal criteria? I see two considerations that weigh against the intuitive approach. First, simplicity is not the only goal—there are often others, such as efficiency, fairness, or revenue. Unless we put simplicity on equal mathematical footing, we cannot study these trade-offs systematically. Without formal criteria for simplicity, we may fixate on other desiderata that are well-formalized.

Second, our intuitive judgements about simplicity are distorted by economic training. Experts in mechanism design suffer from the curse of knowledge; it is hard for them to adopt the perspective of people who do not know game theory (Camerer, Loewenstein, and Weber 1989). In order to produce scientific knowledge about simplicity, we must construct theories that can be tested with data. There is a pressing need for experiments, both to test current theories about simplicity and to discover empirical regularities that future theories might explain. In testing theories, bear in mind that some tasks might be easy if done in isolation, but difficult

when they are part of a broader strategic context. Kagel and Levin (2016) survey experiments on auctions, and Hakimov and Kübler (2021) survey experiments on matching mechanisms.

Much remains to be done. The criteria for simplicity discussed here are not exhaustive; there are other dimensions on which mechanisms can be simple or complex. There could be better ways to formalize the same dimensions, or ways that are more tractable, track human behavior more closely, or have firmer cognitive foundations. Moreover, simplicity comes in degrees, and it would be useful to compare the relative simplicity of different mechanisms, or to find ways to say that one mechanism is “as simple as possible” given other constraints. Nagel and Saitto (2023) recently made progress in this direction.

Finally, another important direction for the simplicity literature is to study how to describe and explain mechanisms, to help participants see for themselves that the mechanism is incentive-compatible (Gonczarowski, Heffetz, and Thomas 2023). Real-world mechanisms, such as spectrum auctions and the National Resident Matching Program, often include detailed advice for participants. The Israeli Psychology Master’s Match even provided participants with a general-audience lecture demonstrating that its mechanism was strategy-proof, but some participants were clearly not convinced (Hassidim, Romm, and Shorrer 2021). Few laboratory experiments have studied such advice, possibly due to concerns about experimenter demand effects (De Quidt, Vesterlund, and Wilson 2019). The experiment conducted by Masuda et al. (2022) offers a rare exception.

■ *I thank Tilman Borgers, Eric Chen, Yunseo Choi, Jiangtao Li, Marek Pycia, and Peter Troyan for valuable comments.*

References

- Abdulkadiroğlu, Atila, Parag A. Pathak, Alvin E. Roth, and Tayfun Sönmez. 2005. “The Boston Public School Match.” *American Economic Review* 95 (2): 368–71.
- Anderson, Ashton, Jon Kleinberg, and Sendhil Mullainathan. 2017. “Assessing Human Error against a Benchmark of Perfection.” *ACM Transactions on Knowledge Discovery from Data (TKDD)* 11 (4): 1–25.
- Arribillaga, R. Pablo, Jordi Massó, and Alejandro Neme. 2020. “On Obvious Strategy-Proofness and Single-Peakedness.” *Journal of Economic Theory* 186: 104992.
- Arribillaga, R. Pablo, Jordi Massó, and Alejandro Neme. 2023. “All Sequential Allotment Rules Are Obviously Strategy-Proof.” *Theoretical Economics* 18 (3): 1023–61.
- Ashlagi, Itai, and Yannai A. Gonczarowski. 2018. “Stable Matching Mechanisms Are Not Obviously Strategy-Proof.” *Journal of Economic Theory* 177: 405–25.
- Athey, Susan, and Jonathan Levin. 2001. “Information and Competition in US Forest Service Timber Auctions.” *Journal of Political Economy* 109 (2): 375–417.
- Bade, Sophie. 2019. “Matching with Single-Peaked Preferences.” *Journal of Economic Theory* 180: 81–99.
- Bade, Sophie, and Yannai A. Gonczarowski. 2017. “Gibbard-Satterthwaite Success Stories and Obvious

- Strategyproofness." In *Proceedings of the 2017 ACM Conference on Economics and Computation*, 565. New York: ACM.
- Bogomolnaia, Anna, and Hervé Moulin.** 2001. "A New Solution to the Random Assignment Problem." *Journal of Economic Theory* 100 (2): 295–328.
- Breitmoser, Yves, and Sebastian Schweighofer-Kodritsch.** 2022. "Obviousness around the Clock." *Experimental Economics* 25 (2): 483–513.
- Bulow, Jeremy, Jonathan Levin, and Paul Milgrom.** 2009. "Winning Play in Spectrum Auctions." NBER Working Paper 14765.
- Börger, Tilman, and Jingtao Li.** 2019. "Strategically Simple Mechanisms." *Econometrica* 87 (6): 2003–35.
- Camerer, Colin, George Loewenstein, and Martin Weber.** 1989. "The Curse of Knowledge in Economic Settings: An Experimental Analysis." *Journal of Political Economy* 97 (5): 1232–54.
- Chatterjee, Kalyan, and William Samuelson.** 1983. "Bargaining under Incomplete Information." *Operations Research* 31 (5): 835–51.
- Cohen, Shani, and Shengwu Li.** 2022. "Sequential Cursed Equilibrium." <https://doi.org/10.48550/arXiv.2212.06025>.
- Dasgupta, Partha, Peter Hammond, and Eric Maskin.** 1979. "The Implementation of Social Choice Rules: Some General Results on Incentive Compatibility." *Review of Economic Studies* 46 (2): 185–216.
- de Quidt, Jonathan, Lise Vesterlund, and Alistair J. Wilson.** 2019. "Experimenter Demand Effects." In *Handbook of Research Methods and Applications in Experimental Economics*, edited by Arthur Schram and Aljaž Ule, 384–400. Cheltenham, UK: Edward Elgar Publishing.
- Dreyfuss, Bnaya, Ori Heffetz, and Matthew Rabin.** 2022. "Expectations-Based Loss Aversion May Help Explain Seemingly Dominated Choices in Strategy-Proof Mechanisms." *American Economic Journal: Microeconomics* 14 (4): 515–55.
- Dubins, Lester E., and David A. Freedman.** 1981. "Machiavelli and the Gale-Shapley Algorithm." *American Mathematical Monthly* 88 (7): 485–94.
- Edelman, Benjamin, Michael Ostrovsky, and Michael Schwarz.** 2007. "Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords." *American Economic Review* 97 (1): 242–59.
- Esponda, Ignacio, and Emanuel Vespa.** 2014. "Hypothetical Thinking and Information Extraction in the Laboratory." *American Economic Journal: Microeconomics* 6 (4): 180–202.
- Esponda, Ignacio, and Emanuel Vespa.** 2023. "Contingent Thinking and the Sure-Thing Principle: Revisiting Classic Anomalies in the Laboratory." *Review of Economic Studies*. <https://doi.org/10.1093/restud/rdad102>.
- Gibbard, Allan.** 1973. "Manipulation of Voting Schemes: A General Result." *Econometrica* 41 (4): 587–601.
- Glazer, Jacob, and Ariel Rubinstein.** 1996. "An Extensive Game as a Guide for Solving a Normal Game." *Journal of Economic Theory* 70 (1): 32–42.
- Golowich, Louis, and Shengwu Li.** 2021. "On the Computational Properties of Obviously Strategy-Proof Mechanisms." <https://doi.org/10.48550/arXiv.2101.05149>.
- Gonczarowski, Yannai A., Ori Heffetz, and Clayton Thomas.** 2023. "Strategyproofness-Exposing Mechanism Descriptions." NBER Working Paper 31506.
- Hakimov, Rustamdjan, and Dorothea Kübler.** 2021. "Experiments on Centralized School Choice and College Admissions: A Survey." *Experimental Economics* 24 (2): 434–88.
- Harstad, Ronald M.** 2000. "Dominant Strategy Adoption and Bidders' Experience with Pricing Rules." *Experimental Economics* 3: 261–80.
- Hassidim, Avinatan, Assaf Romm, and Ran I. Shorrer.** 2021. "The Limits of Incentives in Economic Matching Procedures." *Management Science* 67 (2) 951–63.
- Hurwicz, Leonid.** 1973. "The Design of Mechanisms for Resource Allocation." *American Economic Review* 63 (2) 1–30.
- Jéhiel, Philippe.** 1995. "Limited Horizon Forecast in Repeated Alternate Games." *Journal of Economic Theory* 67 (2): 497–519.
- Jéhiel, Philippe, and Dov Samet.** 2007. "Valuation Equilibrium." *Theoretical Economics* 2 (2): 163–85.
- Kagel, John H.** 2016. "Auctions: A Survey of Experimental Research." In *The Handbook of Experimental Economics*, edited by John H. Kagel and Alvin E. Roth, 501–86. Princeton: Princeton University Press.
- Kagel, John H., Ronald M. Harstad, and Dan Levin.** 1987. "Information Impact and Allocation Rules in Auctions with Affiliated Private Values: A Laboratory Study." *Econometrica* 55 (6): 1275–1304.

- Kagel, John H., and Dan Levin.** 1993. "Independent Private Value Auctions: Bidder Behaviour in First-, Second-, and Third-Price Auctions with Varying Numbers of Bidders." *Economic Journal* 103 (419): 868–79.
- Karp, Richard M.** 1972. "Reducibility among Combinatorial Problems." In *Complexity of Computer Computations*, edited by Raymond E. Miller and James W. Thatcher, 85–103. New York: Plenum Press.
- Klemperer, Paul.** 2010. "The Product-Mix Auction: A New Auction Design for Differentiated Goods." *Journal of the European Economic Association* 8 (2-3): 526–36.
- Leyton-Brown, Kevin, Paul Milgrom, and Ilya Segal.** 2017. "Economics and Computer Science of a Radio Spectrum Reallocation." *Proceedings of the National Academy of Sciences* 114 (28): 7202–09.
- Li, Jiangtao, and Piotr Dworczak.** 2024. "Are Simple Mechanisms Optimal When Agents Are Unsophisticated?" Unpublished.
- Li, Shengwu.** 2017. "Obviously Strategy-Proof Mechanisms." *American Economic Review* 107 (11): 3257–87.
- Li, Shengwu, and Irene Lo.** 2023. "New Solution Concepts." In *Online and Matching-Based Market Design*, edited by Federico Echenique, Nicole Immorlica, and Vijay V. Vazirani, 575–90. Cambridge, UK: Cambridge University Press.
- Lucier, Brendan.** 2017. "An Economic View of Prophet Inequalities." *ACM SIGecom Exchanges* 16 (1): 24–47.
- Lucking-Reiley, David.** 2000. "Vickrey Auctions in Practice: From Nineteenth-Century Philately to Twenty-First-Century E-commerce." *Journal of Economic Perspectives* 14 (3): 183–92.
- Mackenzie, Andrew.** 2020. "A Revelation Principle for Obviously Strategy-Proof Implementation." *Games and Economic Behavior* 124: 512–33.
- Mandal, Pinaki, and Souvik Roy.** 2022. "Obviously Strategy-Proof Implementation of Assignment Rules: A New Characterization." *International Economic Review* 63 (1): 261–90.
- Marszalec, Daniel, Alexander Teytelboym, and Sanna Laksá.** 2020. "Epic Fail: How Below-Bid Pricing Backfires in Multiunit Auctions." CIRJE Working Paper CIRJE-F-1096.
- Martínez-Marquina, Alejandro, Muriel Niederle, and Emanuel Vespa.** 2019. "Failures in Contingent Reasoning: The Role of Uncertainty." *American Economic Review* 109 (10): 3437–74.
- Masuda, Takehito, Ryo Mikami, Toyotaka Sakai, Shigehiro Serizawa, and Takuma Wakayama.** 2022. "The Net Effect of Advice on Strategy-Proof Mechanisms: An Experiment for the Vickrey Auction." *Experimental Economics* 25 (3): 902–41.
- McCabe, Kevin A., Stephen J. Rassenti, and Vernon L. Smith.** 1990. "Auction Institutional Design: Theory and Behavior of Simultaneous Multiple-Unit Generalizations of the Dutch and English Auctions." *American Economic Review* 80 (5): 1276–83.
- Meisner, Vincent, and Jonas von Wangenheim.** 2023. "Loss Aversion in Strategy-Proof School-Choice Mechanisms." *Journal of Economic Theory* 207: 105588.
- Milgrom, Paul.** 2000. "Putting Auction Theory to Work: The Simultaneous Ascending Auction." *Journal of Political Economy* 108 (2): 245–72.
- Milgrom, Paul, and Ilya Segal.** 2020. "Clock Auctions and Radio Spectrum Reallocation." *Journal of Political Economy* 128 (1): 1–31.
- Myerson, Roger B.** 1979. "Incentive Compatibility and the Bargaining Problem." *Econometrica* 47 (1): 61–73.
- Nagel, Lea, and Roberto Saitto.** 2023. "A Measure of Complexity for Strategy-Proof Mechanisms." In *Proceedings of the 24th ACM Conference on Economics and Computation*, 1017. New York: ACM.
- Niederle, Muriel, and Emanuel Vespa.** 2023. "Cognitive Limitations: Failures of Contingent Thinking." *Annual Review of Economics* 15: 307–28.
- Pathak, Parag A., and Tayfun Sönmez.** 2008. "Leveling the Playing Field: Sincere and Sophisticated Players in the Boston Mechanism." *American Economic Review* 98 (4): 1636–52.
- Prendergast, Canice.** 2017. "How Food Banks Use Markets to Feed the Poor." *Journal of Economic Perspectives* 31 (4): 145–62.
- Prendergast, Canice.** 2022. "The Allocation of Food to Food Banks." *Journal of Political Economy* 130 (8): 1993–2017.
- Pycia, Marek, and Peter Troyan.** 2023. "A Theory of Simplicity in Games and Mechanism Design." *Econometrica* 91 (4): 1495–526.
- Rees-Jones, Alex.** 2018. "Suboptimal Behavior in Strategy-Proof Mechanisms: Evidence from the Residency Match." *Games and Economic Behavior* 108: 317–30.
- Rees-Jones, Alex, and Samuel Skowronek.** 2018. "An Experimental Investigation of Preference Misrepresentation in the Residency Match." *Proceedings of the National Academy of Sciences* 115 (45): 11471–76.

- Roth, Alvin E.** 2002. "The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics." *Econometrica* 70 (4): 1341–78.
- Roth, Alvin E., and Elliott Peranson.** 1999. "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design." *American Economic Review* 89 (4): 748–80.
- Shorrer, Ran I., and Sándor Sóvágó.** 2023. "Dominated Choices in a Strategically Simple College Admissions Environment." *Journal of Political Economy Microeconomics* 1 (4): 781–807.
- Thomas, Clayton.** 2020. "Classification of Priorities Such That Deferred Acceptance is Obviously Strategyproof." <https://doi.org/10.48550/arXiv.2011.12367>.
- Troyan, Peter.** 2019. "Obviously Strategy-Proof Implementation of Top Trading Cycles." *International Economic Review* 60 (3): 1249–61.
- Vickrey, William.** 1961. "Counterspeculation, Auctions, and Competitive Sealed Tenders." *Journal of Finance* 16 (1): 8–37.
- Williams, Kevin Jon.** 1995. "A Reexamination of the NRMP Matching Algorithm. National Resident Matching Program." *Academic Medicine: Journal of the Association of American Medical Colleges* 70 (6): 470–76.
- Wilson, Robert.** 1987. "Game-Theoretic Analyses of Trading Processes." In *Advances in Economic Theory*, edited by Truman Fassett Bewley, 33–70. Cambridge, UK: Cambridge University Press.
- Wilson, Robert.** 2002. "Architecture of Power Markets." *Econometrica* 70 (4): 1299–340.

