

Recognizing Dominance through Local Improvement*

Job Market Paper

[Click here for the latest version](#)

Linchen Zhang[†]

October 29, 2026[‡]

Abstract

Strategy-proofness is the traditional standard for predictable behavior, yet participants often fail to play dominant strategies. Most existing simplicity concepts attribute these failures to the cost of contingent reasoning, which they lower by comparing the worst case of the dominant strategy with the best case of an alternative. I argue that there is another approach, one that makes more use of the structure of the strategy set, which I formalize as a neighbor relation. Comparing neighboring strategies can reduce the need for contingent reasoning, and, even with a single contingency, it can simplify the search through the strategy set. I introduce a new simplicity criterion: a mechanism is local improvement solvable under a neighbor relation if repeatedly moving to an improving neighbor always reaches a dominant strategy. Strategy-proofness is the criterion under the coarsest relation, and many classical mechanisms satisfy it under finer ones. In extensive-form mechanisms without duplicate dominant strategies, the criterion under one-situation neighbors is strict action-proofness, which is distinct from obvious strategy-proofness. Among local improvement solvable mechanisms that implement the same outcomes, the criterion predicts more dominant-strategy play where the local comparison is salient.

JEL codes: C72, C91, D44, D47, D82, D91.

*I am deeply indebted to Songzi Du for his invaluable guidance on this project. I am immensely grateful to Aram Grigoryan, Joel Sobel, Joel Watson, and Emanuel Vespa for their useful suggestions for this paper. I also thank Isabel Trevino, Denis Shishkin, Vince Crawford, Suraj Malladi, and participants at the UCSD TBE seminar and workshop for valuable discussions. Claude was used for text refinement and proof verification in this project. All errors are my own.

[†]Department of Economics, UC San Diego. Email: liz032@ucsd.edu. Website: linchenzhang.github.io.

[‡]This is a preliminary version. It may be incorporated into a new version with additional experimental results from the experiment proposed in Appendix C.4, which I plan to run in the near future.

1 Introduction

A mechanism can be relied on to deliver the outcomes it was designed for only if participants behave as predicted. The traditional standard for predictable behavior is strategy-proofness. In a strategy-proof mechanism, every participant has a dominant strategy, one that is optimal whatever the others do. The prediction then does not depend on her beliefs about the others, and it should be relatively simple for the participant to figure out what her dominant strategy is. Yet participants often fail to play their dominant strategies. [Kagel, Harstad and Levin \(1987\)](#) find that subjects overbid in the sealed-bid second-price auction, and [Rees-Jones \(2018\)](#) and [Hassidim, Romm and Shorrer \(2021\)](#) find that applicants misreport in high-stakes deferred acceptance clearinghouses. The rate of dominant-strategy play also differs widely across strategy-proof mechanisms, as [Cason et al. \(2006\)](#) and [Chen and Sönmez \(2006\)](#) show. Some strategy-proof mechanisms therefore seem to be simpler than others, and strategy-proofness by itself does not say which.

One reason a dominant strategy can be hard to recognize is that recognizing it may require heavy contingent reasoning.¹ Suppose you value an object at \$3.50 and bid for it in a sealed-bid second-price auction. Bids are whole dollars from \$0 to \$7, and a tie for the highest bid means no sale. In this setting, your dominant bid is \$4, the lowest bid above your value. To compare \$4 directly with a distant alternative such as \$7, you must reason contingency by contingency. The two bids give you different payoffs only when the highest rival bid is \$4, \$5, or \$6. In each of these cases, \$7 wins at a price above your value, while \$4 does not win. To conclude that \$4 is better than \$7, you must therefore work through three cases, which participants can find burdensome.

In response to failures of contingent reasoning, the recent literature on simple mechanisms looks for mechanisms that require less contingent reasoning. The most influential concept in this literature is the obvious strategy-proofness of [Li \(2017\)](#). A strategy is obviously dominant if, at every situation it reaches, its worst possible continuation outcome is at least as good as the best possible outcome of any deviation at that situation. Here a situation is a point at which the player must choose an action. A mechanism is obviously strategy-proof, or OSP, when every type of every player has an obviously dominant strategy. A participant can then recognize dominance by this worst-best comparison, without separating the contingencies. The sealed-bid second-price auction is not OSP, since the best that \$7 can bring, winning at a price of \$0, beats the worst that \$4 can bring, winning nothing. Later concepts, reviewed in Section 2, strengthen or weaken obvious strategy-proofness, or rank strategy-proof mechanisms by the contingent reasoning that remains under worst-best comparisons, as [Nagel and Saitto \(forthcoming\)](#) do. I

¹[Niederle and Vespa \(2023\)](#) review the experimental evidence on failures of contingent reasoning.

refer to obvious strategy-proofness and the later concepts that compare a worst case with a best case as the worst-best comparison concepts.

Worst-best comparisons may not, however, be the only simplifying route that participants can use to recognize a dominant strategy. In this paper, I study a second way to reduce contingent reasoning, the comparison of neighboring strategies. The idea is that two similar strategies, which I formally define as neighbors, lead to the same outcome, or call for the same reasoning, in many contingencies. When a participant compares two neighbors, the contingencies in which they lead to the same outcome cancel out, and less contingent reasoning is needed. I argue that the comparison of neighbors defines a dimension of simplicity that the worst-best comparison concepts do not measure.

For example, return to the sealed-bid second-price auction and let the neighbors of a bid be the bids one dollar above and below it. The bids \$4 and \$3 are neighbors and give you different payoffs only when the highest rival bid is exactly \$3. Then \$4 wins at a price of \$3, worth \$0.50 to you, and \$3 does not win. So \$4 improves \$3, in the sense that \$4 is at least as good whatever the rivals bid and strictly better in at least one case. Similarly, \$4 improves \$5. The two bids give different payoffs only when the highest rival bid is \$4, and then \$5 wins at a price above your value, which gives you a negative payoff. By the same argument, each bid other than \$4 is improved by its neighbor closer to \$4. Moving from any bid to an improving neighbor, as long as one exists, therefore ends at \$4, the dominant bid.

The argument with neighboring bids uses an aspect of the mechanism that the worst-best comparison concepts do not fully exploit: the structure of the strategy set. The strategy set has an order, here the order of bids, that determines which strategies are neighbors. As the auction example shows, the order of strategies can be a further source of simplicity in recognizing the dominant strategy. To formalize the argument, I add one primitive to the standard model, a neighbor relation on each strategy set. The neighbor relation specifies which pairs of strategies are neighbors, and so which pairs are compared locally. A relation of this kind is implicit in the one-shot deviation principle of game theory. The principle, presented for instance in [Osborne and Rubinstein \(1994\)](#), checks a strategy only against the strategies that differ from it at one situation. A neighbor relation allows other notions of a single change, such as one grid step of a bid or an exchange of two adjacent entries in a ranking.

A strategy improves another if it is at least as good in every contingency and strictly better in some. The local improvement algorithm moves to an improving neighbor as long as one exists. I then propose the key simplicity criterion of this paper, local improvement solvability. A mechanism is local improvement solvable under a neighbor relation if, for every type of every player, every run of the algorithm ends at a dominant strategy, whatever the starting strategy. Local improvement solvability implies strategy-proofness and becomes stronger as

the neighbor relation becomes finer, that is, as each strategy has fewer neighbors. Under the complete neighbor relation, in which every two strategies are neighbors, local improvement solvability is exactly strategy-proofness, which is therefore the coarsest member of the resulting family of conditions.

I then show that many classical strategy-proof mechanisms satisfy the criterion under their natural neighbor relations, the relations that the presentations of their strategy sets define. The sealed-bid second-price auction and the multi-unit Vickrey auction are local improvement solvable under one-step neighbors, which change a bid, or one marginal bid, by one grid step. In matching mechanisms in which participants submit rankings, two rankings are one-exchange neighbors if they differ by an exchange of two adjacent entries. On the full domain of rankings, deterministic matching mechanisms of this kind in which truthful reporting is dominant are local improvement solvable under one-exchange neighbors, provided the environment is rich. Richness, which seems a mild condition, requires every exchange of two adjacent entries to change the participant's outcome in some contingency.² Serial dictatorship, top trading cycles, and deferred acceptance are such mechanisms. In these classical mechanisms, the comparison between two neighboring strategies involves only one case. They may therefore contain a hidden, shorter path to the dominant strategy, and Section 7 discusses how a mechanism can guide participants along it.

Local improvement solvability extends to extensive-form mechanisms and games. At a situation h , the extensive-form structure restricts the comparison to the strategies and contingencies that reach h . The reduced normal-form mechanism does not show this restriction, which the worst-best comparison concepts use. Local improvement solvability can also use a finer relation. Under one-situation neighbors, two strategies are neighbors if they differ in the action at exactly one situation. A step of the local improvement algorithm then changes one action and holds the rest of the strategy fixed. Such a step is a one-shot deviation, so one-situation neighbors are the relation implicit in the one-shot deviation principle.

Action-proofness, a condition I introduce, concerns exactly these one-shot deviations.³ It requires that, at each situation the recommended strategy reaches, no one-shot deviation from the recommended action pay more, whatever the player does at her other situations and whatever the others do. Theorem 2 shows that local improvement solvability under one-situation neighbors is equivalent to strict action-proofness, absent duplicate dominant strategies. The theorem can thus be read as a one-shot deviation principle for dominance. In a strictly

²Richness fails, for example, when the participant knows that some adjacent pair on her list cannot affect her outcome.

³Action-proofness has precedents in the agent normal form of [Selten \(1975\)](#) and in the weak dominance between actions of [Catonini and Xue \(2026\)](#).

action-proof mechanism, every strategy with no improving one-shot deviation is dominant. Under action-proofness the participant need not plan her own later actions, but she must still compare actions contingency by contingency, which obvious strategy-proofness lets her avoid. Neither condition implies the other. The strong obvious strategy-proofness of [Pycia and Troyan \(2023\)](#), or SOSPP, implies both action-proofness and obvious strategy-proofness, although the two together do not imply SOSPP.

The worst-best comparison concepts and local improvement solvability can disagree, as three example games in Section 6 show. In the first example, the pick-a-number game, two players each pick a number from $\{0, 1, \dots, K\}$, and each receives her own number plus λ times the other player's number, where $\lambda > K$. Raising one's number by one raises one's payoff by exactly one, whatever the other player picks, so the game is local improvement solvable under one-step neighbors and seems easy. Yet picking K is not obviously dominant, since its worst payoff, K , falls short of the best payoff $a + \lambda K$ of any smaller number a . The measure of [Nagel and Saitto \(forthcoming\)](#) classifies the game as maximally complex, and no strategy-proof extensive-form game that implements the same outcomes is OSP. The worst-best comparison concepts seem to mislabel this easy game as hard, while local comparisons explain its simplicity. Chess, when White has a winning strategy, and Sudoku with a unique solution show the opposite error. Both games seem difficult, yet chess is OSP for White and Sudoku is SOSPP, so both are simple by obvious strategy-proofness. On the other hand, neither game is local improvement solvable under its natural neighbor relation, one-situation neighbors in chess and one-cell neighbors in Sudoku. In each game, some strategy with two mistakes has no improving neighbor, although that strategy is not dominant. Local improvement solvability therefore accounts for part of what the worst-best comparison concepts miss.

Sudoku also shows that local improvement solvability concerns more than contingent reasoning. A Sudoku puzzle has a single contingency, since there is no other player and no chance move, so finding the solution requires no contingent reasoning at all. Its difficulty seems to lie instead in the search for the solution among the many fillings of the grid. With a single contingency, local improvement solvability requires every strategy without an improving neighbor to be a best strategy, as in an optimization problem in which every local maximum is global. The neighbor relation thus captures a source of simplicity beyond contingent reasoning. It concerns the search through the strategy set, and the worst-best comparison concepts do not measure it. The example in Section 3.2 that separates local improvement solvability from strategy-proofness makes the same point with three strategies and a single contingency.

Local improvement solvability also yields a behavioral prediction. Comparisons between neighbors help a participant only if she makes them. A mechanism makes the local comparison salient when it asks her for a choice that contains a comparison between two neighboring

strategies, or when it resolves a submitted strategy step by step. A request for her whole strategy at once does neither. I therefore predict that, of two strategy-proof mechanisms that implement the same outcomes and are local improvement solvable for the same reason, the one that makes the local comparison salient produces more dominant-strategy play.

The ascending clock auction is OSP and produces the same outcomes as the sealed-bid second-price auction. [Li \(2017\)](#) finds more dominant-strategy play in the ascending clock auction and attributes the difference to obvious strategy-proofness. Yet the ascending clock auction also asks at each price whether to exit now or stay one step longer, a comparison of two neighboring bids. Obvious strategy-proofness and the salience of the local comparison are thus two separate channels through which the ascending clock auction could help the participant recognize the dominant strategy. Li's comparison between the two auctions changes both channels at once.

A later experiment separates the two channels. [Breitmoser and Schweighofer-Kodritsch \(2022\)](#) add to the design of [Li \(2017\)](#) a passive clock treatment, in which subjects submit a sealed bid and then watch an ascending clock, driven by that bid, determine the outcome. The passive clock has the same extensive-form mechanism as the sealed-bid second-price auction and is not OSP, so obvious strategy-proofness on its own predicts no difference between the two. Yet it makes the local comparison salient. It removes about half of the difference in deviations from truthful bidding between the sealed-bid second-price auction and the ascending clock auction in which bidders see their rivals exit. Dynamic bidding, the feature that makes the ascending clock auction OSP, accounts for a further twelve percent, which is not statistically significant. Showing bidders when their rivals exit accounts for almost forty percent, which I attribute to learning from the environment. The results suggest that, in the ascending clock auction, the local comparison does more than obvious strategy-proofness to help participants recognize the dominant strategy.⁴

Table [B.1](#) in Appendix [B](#) collects further experimental evidence on the effectiveness of obvious strategy-proofness and of the local comparison. The evidence seems favorable to both channels on balance. Each appears to help participants recognize the dominant strategy, substantially in some environments and marginally in others, although in some settings the gain does not carry over to outcomes.

The local improvement criterion leads to two design recommendations. The first is to choose, among mechanisms that implement the same outcomes, one that is local improvement solvable under its natural neighbor relation. The second is to make the local comparison salient. Local improvement solvability resembles a short solution to a difficult mathematics problem,

⁴Section [7.4](#) discusses another explanation of the passive clock's gain, learning from feedback.

which a student may miss on her own but find with the right hint. A presentation that makes the local comparison salient is a hint of this kind. If the salience channel is real, the hint is also cheap, since a presentation that only makes the local comparison salient leaves the incentive properties of the mechanism unchanged.

The rest of the paper is organized as follows. Section 2 reviews the related literature, including the closest antecedent, [Catonini and Xue \(2026\)](#). Sections 3 to 5 define the criterion and apply it to classical and extensive-form mechanisms. Section 6 examines the three games, Section 7 states the prediction and reviews the evidence, and Section 8 gives the design recommendations and concludes. Proofs omitted from the main text are in Appendix A, Appendix B collects the experimental evidence, and Appendix C contains further discussion, including a proposed design.

2 Related literature

The literature on the worst-best comparison concepts begins with the obvious strategy-proofness of [Li \(2017\)](#). The concepts that followed vary Li’s comparison along two dimensions, and Table 1 in Section 5.3 makes both explicit. The first dimension is how the contingencies are compared. Dominance compares the recommended strategy with a deviation one contingency at a time. Obvious dominance compares the worst case of the recommended strategy against the best case of the deviation. [Zhang and Levin \(2017\)](#) give a decision-theoretic foundation for obvious dominance, and the concepts of [Zhang and Levin \(2018\)](#) and [Arribillaga, Massó and Neme \(2024\)](#) lie between the two. They compare the worst case against the best case within each element of a partition of the contingencies, which [Arribillaga, Massó and Neme \(2024\)](#) derive from a partition of the players. [Chew and Wang \(2022\)](#) weaken obvious dominance using the sure-thing principle. The second dimension is how much of her own future play the participant is assumed to plan. [Pycia and Troyan \(2023\)](#) define strong obvious strategy-proofness, under which she plans none of it. They also define a family of concepts indexed by her planning horizon, the set of future decisions she plans for. In this family, obvious strategy-proofness is the case in which she plans for all of them. [Nagel and Saitto \(forthcoming\)](#) turn the binary classification into a complete ranking of strategy-proof mechanisms by the contingent reasoning they require.

Two further concepts concern what a participant notices or needs to believe rather than what she must compare. The obvious manipulations of [Troyan and Morrill \(2020\)](#) are the deviations a participant can see to be profitable when she compares worst cases with worst cases and best cases with best cases. In the strategically simple mechanisms of [Börger and Li \(2019\)](#),

she can determine an optimal choice from first-order beliefs about the others' preferences and first-order certainty of their rationality, without higher-order beliefs.⁵ None of these concepts uses the structure of the strategy set beyond what an extensive-form mechanism records.

This paper adds the structure of the strategy set as an input and replaces the comparison of the recommended strategy with each alternative by a chain of comparisons between neighbors. The two kinds of argument are complementary. Section 5 shows that strategy-proofness, OSP, and SOSp are conditions on a single step of the local improvement algorithm in extensive-form mechanisms. Local improvement solvability instead requires every run of the algorithm to end at a dominant strategy.

Catonini and Xue (2026) is the closest antecedent. They compare two actions at one situation rather than two strategies. Given a partition of the contingencies that reach the situation, one action dominates another if every continuation after the second is matched by some continuation after the first. In each element of the partition, the worst payoff of the match must be at least the best payoff of the continuation it matches. Their weak dominance between actions uses the singleton partition. They show that a strategy is dominant if and only if each action it prescribes at a situation it reaches is weakly dominant. The trivial partition gives the corresponding result for obvious dominance.

Their weak dominance between actions is closely related to action dominance, which Section 5 introduces and Theorem 2 relates to local improvement solvability under one-situation neighbors. Both compare two actions at one situation, one contingency at a time. Their condition lets the continuation after the dominating action depend on the continuation it is matched with. Action dominance requires every continuation after the recommended action to do at least as well, in every contingency, as every continuation after the other action. Catonini and Xue (2026) do not study this stronger condition. Section 5.3 shows that this requirement is exactly what separates action-proofness from strategy-proofness.

The local dominance of Catonini and Xue (2026) derives the partition from the game. Where the dominating action ends the player's play, its worst payoff must be at least the best payoff of the other action. Elsewhere, "continuing in the same way" after the two actions, which repeats after the dominating action the later actions taken after the other one, must give the same outcome. A mechanism is locally strategy-proof if each type has a locally dominant action at every situation she reaches by choosing such actions, and local strategy-proofness implies strategy-proofness. Local dominance is close to the comparison between neighbors. In the English auction with alternating bids, their argument compares leaving now with bidding and leaving at the next turn. Section 7.1 identifies the same comparison, between adjacent

⁵Li (2024) reviews three formalizations of simplicity.

exit prices, in the ascending clock auction. Like the natural neighbor relation of this paper, local dominance depends on how the game is presented, since “continuing in the same way” is defined by stages and action labels. Appendix C.2 relates the conditions of [Catonini and Xue \(2026\)](#) to those of Table 1.

This paper differs from theirs in three respects. First, the comparison is between two strategies that are neighbors under a relation on the strategy set, and such relations are defined for static mechanisms as well. Local strategy-proofness coincides with obvious strategy-proofness in static mechanisms, as they show, so the sealed-bid second-price auction, which is local improvement solvable under one-step neighbors, is not locally strategy-proof. Second, a local improvement is usually judged one contingency at a time, while local dominance compares a worst case with a best case where the dominating action ends the player’s play. Third, the object of study is whether repeated local comparisons reach the dominant strategy from every starting strategy. The resulting family of conditions is ordered by the fineness of the neighbor relation, with strategy-proofness as its coarsest member.

The local incentive compatibility literature studies when local incentive constraints imply global ones. [Carroll \(2012\)](#) studies the question for cardinal and ordinal preferences and with transfers. On ordinal domains the question is also studied in [Sato \(2013\)](#), [Cho \(2016\)](#), and [Kumar et al. \(2021\)](#). With transfers it is studied in [Archer and Kleinberg \(2014\)](#), [Mishra, Pramanik and Roy \(2016\)](#), and [Kumar and Roy \(2024\)](#), and for random assignment in [Mennle and Seuken \(2021\)](#). These papers ask whether an analyst who knows the outcome function can verify strategy-proofness from the local constraints alone. This paper asks whether a participant who compares neighboring strategies reaches the dominant strategy.⁶

[Glazer and Rubinstein \(1996\)](#) show that a static game solvable by iterated elimination of dominated strategies can be presented as an extensive-form game with the same normal-form game. In the extensive-form game, backward induction reaches the same solution through the same comparisons. They argue that the presentation guides the player to the solution. Section 7 makes the corresponding point for strategy-proof mechanisms, arguing that a dynamic mechanism or a passive presentation guides the participant by making the local comparison salient. Table B.1 in Appendix B collects the experiments that hold the outcomes fixed and vary the mechanism or its presentation.⁷

⁶Proposition 4(b) shows that in a deterministic ranking mechanism on the full strict domain that is strategy-proof under truthful reporting, an adjacent exchange changes the outcome only between the two exchanged objects. For deterministic mechanisms, this result is one direction of the characterization of strategy-proofness by swap monotonicity and upper and lower invariance in [Mennle and Seuken \(2021\)](#). It also restates the mistake monotonicity of [Cho \(2016\)](#) for this environment. Part (b) is the direction from global to local constraints.

⁷Among these experiments, [Gonczarowski, Heffetz and Thomas \(2023\)](#) also formalize descriptions of a mechanism as directed graphs that compute its outcome. [Kagel and Levin \(2001\)](#) report related experiments on multi-unit auctions, and [Hakimov and Kübler \(2021\)](#) survey the matching experiments.

Two further lines of literature concern the errors that remain. [Babaioff, Dobzinski and Kupfer \(2023\)](#) define the primitive complexity of a menu as the number of operations needed to find its best element. [Oprea \(2020\)](#) and [Banovetz and Oprea \(2023\)](#) find that the number of states in a rule predicts its cost and use. [Enke and Shubatt \(2023\)](#) and [Oprea \(2024\)](#) relate errors in choices between lotteries to their complexity. Primitive complexity is the kind of difficulty the Sudoku game of Section 6 has, and the worst-best comparison concepts classify it as absent. [Velez and Brown \(2022\)](#) ask which deviations from the dominant strategy are consistent with empirical equilibrium, a refinement of Nash equilibrium.⁸ Their criterion asks which deviations survive as limits of noisy behavior in which an action is played more often than another only if it pays more. Local improvement solvability concerns which deviations survive the participant's own comparison of neighboring strategies.

3 The criterion

3.1 Basic setup

A mechanism Γ specifies for each player i a finite strategy set S_i and an outcome function $f : S \rightarrow \Omega$ that maps strategy profiles $s = (s_0, s_1, \dots, s_n) \in S = S_0 \times S_1 \times \dots \times S_n$ to outcomes, where $s_0 \in S_0$ is nature's move. Player i has a type θ_i in a set Θ_i and a utility $u_i(s_i, s_{-i}; \theta_i)$ over the outcome of each profile. Throughout, the type is carried explicitly as an argument of the utility, and s_{-i} ranges over the contingencies, the profiles of the others' strategies and nature's moves.

The new primitive is a neighbor relation on each strategy set.

Definition 1 (Neighbor relation): *A neighbor relation on S_i assigns to each strategy s_i a set $N_i(s_i) \subseteq S_i \setminus \{s_i\}$ of neighbors, with $s'_i \in N_i(s_i)$ if and only if $s_i \in N_i(s'_i)$.*

Definition 2 (Finer and complete neighbor relations): *A neighbor relation N_i is finer than N'_i if $N_i(s_i) \subseteq N'_i(s_i)$ for every s_i . The complete neighbor relation is $N_i(s_i) = S_i \setminus \{s_i\}$.*

Three relations on bids illustrate the definitions. Let a strategy be a bid on the grid $S_i = \{0, 1, \dots, K\}$. Under one-step neighbors, the neighbors of b are $b - 1$ and $b + 1$. Under two-step neighbors, they are the bids at most two steps from b , namely $b - 2, b - 1, b + 1$, and $b + 2$. Under the complete neighbor relation, they are all other bids. In each case only bids on

⁸[Velez and Brown \(2022\)](#) find that deviations that change the outcome are consistent with empirical equilibrium only when the rule is bossy and information is not interior.

the grid count. One-step neighbors are finer than two-step neighbors, and both are finer than the complete neighbor relation, as every neighbor relation is.

The choice of neighbor relation matters, since a mechanism can be local improvement solvable under one relation and not under another. Section 7 argues that a dynamic mechanism can determine which neighbors a participant considers.

Definition 3 (Dominance and strategy-proofness): A strategy s_i^* is dominant for θ_i if

$$u_i(s_i^*, s_{-i}; \theta_i) \geq u_i(s_i, s_{-i}; \theta_i) \quad \text{for all } s_i \in S_i \text{ and all } s_{-i} \in S_{-i}. \quad (1)$$

The mechanism Γ is strategy-proof (SP) if every type of every player has a dominant strategy.

Two strategies are payoff-equivalent for a type if they give her the same utility in every contingency. Two dominant strategies must be payoff-equivalent, since each is at least as good as the other in every contingency.

Definition 4 (Recommendation): A recommendation is a profile of maps $\sigma_i : \Theta_i \rightarrow S_i$. The mechanism Γ is strategy-proof under σ if each $\sigma_i(\theta_i)$ is dominant for θ_i .

In a direct mechanism, where $S_i = \Theta_i$, the recommendation is truthful reporting, $\sigma_i(\theta_i) = \theta_i$. In the ascending clock auction of Section 7.1 it is to exit at the first price at or above the value. In the dynamic serial dictatorship of Section 7.3 it is to take the top available object. The results below concern whether comparisons between neighboring strategies lead to $\sigma_i(\theta_i)$. They do not use that $\sigma_i(\theta_i)$ is a reported type.

Definition 5 (Improvement): A strategy s'_i improves s_i for θ_i if $u_i(s'_i, s_{-i}; \theta_i) \geq u_i(s_i, s_{-i}; \theta_i)$ for every $s_{-i} \in S_{-i}$, with strict inequality for some $s_{-i} \in S_{-i}$.

No strategy improves itself, and no strategy improves a dominant strategy. If s''_i improves s'_i and s'_i improves s_i , then s''_i improves s_i .

3.2 Local improvement solvability

Definition 6 (Local improvement): A strategy s'_i locally improves s_i for θ_i if $s'_i \in N_i(s_i)$ and s'_i improves s_i . Write $s_i \rightarrow s'_i$.

In words, a local improvement is a neighbor that improves the current strategy, one at least as good in every contingency and strictly better in some. Compare the local improvement with (1). The local comparison is the same pointwise comparison that strategy-proofness makes, restricted to a pair of neighbors, so it does not require beliefs about the others. In the

applications of Section 4, the local comparison reduces to one case, the contingencies in which the two neighbors give the participant different outcomes.

Definition 7 (Local improvement algorithm): *The local improvement algorithm for θ_i starts at some $s_i \in S_i$ and, while the current strategy has at least one local improvement, moves to one of them, chosen arbitrarily. A run is the sequence of strategies visited, $s_i^0 \rightarrow s_i^1 \rightarrow \dots \rightarrow s_i^k$. A strategy with no local improvement is locally unimprovable.*

Lemma 1: *Under any neighbor relation, every run of the local improvement algorithm is finite and ends at a locally unimprovable strategy.*

Proof. Each strategy in a run improves the one before it. The improvement relation is transitive, so if a run returned to a strategy it had visited, that strategy would improve itself, which is impossible. Hence no run visits a strategy twice, every run has at most $|S_i| - 1$ steps, and the last strategy of a run has no local improvement. $\square$

Definition 8 (Local improvement solvability): *The mechanism Γ is local improvement solvable for θ_i under N_i if all locally unimprovable strategies of θ_i are payoff-equivalent. It is local improvement solvable under $N = (N_1, \dots, N_n)$ if it is local improvement solvable for every type θ_i of every player i under N_i .*

Under the definition, every run of the algorithm stops at the same strategy up to payoff equivalence, from any starting strategy and with any choice of improvements along the way.⁹

Theorem 1: (a) *If Γ is local improvement solvable for θ_i under a neighbor relation N_i , its locally unimprovable strategies under N_i are exactly the dominant strategies of θ_i . Hence local improvement solvability under any neighbor relation implies strategy-proofness.*

(b) *Strategy-proofness does not imply local improvement solvability under one-step neighbors.*

Proof. Part (a) follows from transitivity. Fix any strategy s_i . By Lemma 1 a run from s_i ends at a locally unimprovable strategy p , and each step of the run is a pointwise inequality, so $u_i(p, s_{-i}; \theta_i) \geq u_i(s_i, s_{-i}; \theta_i)$ for every s_{-i} . Any locally unimprovable strategy p' is payoff-equivalent to p , so $u_i(p', s_{-i}; \theta_i) \geq u_i(s_i, s_{-i}; \theta_i)$ for every s_{-i} . Since s_i was arbitrary, every

⁹The definition does not require the locally unimprovable strategy to be unique. A strategy that is payoff-equivalent to a dominant strategy is itself dominant and therefore locally unimprovable, so uniqueness would rule out every mechanism in which some type has two dominant strategies.

locally unimprovable strategy is dominant. Conversely, no strategy improves a dominant strategy, so a dominant strategy has no local improvement and is locally unimprovable.

Part (b) uses the following counterexample. A single agent with one type chooses a strategy from $\{0, 1, 2\}$. There are no other players and no moves of nature, so there is a single contingency. Strategy 0 pays her 1, strategy 1 pays 0, and strategy 2 pays 2. One strategy thus improves another exactly when it pays more. Under one-step neighbors, as for bids in Section 3.1, the only neighbor of 0 is 1, the neighbors of 1 are 0 and 2, and the only neighbor of 2 is 1.

Strategy 2 pays the most, so it is dominant and the mechanism is strategy-proof. No strategy improves a dominant strategy, so strategy 2 is locally unimprovable. Strategy 0 is locally unimprovable as well, since its only neighbor, strategy 1, pays less. The locally unimprovable strategies 0 and 2 are not payoff-equivalent, since they pay 1 and 2. Hence the mechanism is not local improvement solvable under one-step neighbors. $\square$

The example in part (b) shows the failure that local reasoning allows. A participant who compares strategy 0 against its neighbor finds no improvement, although strategy 2, two steps away, improves it.

The example has a single contingency, so the failure does not come from contingent reasoning. Strategy 2 is even obviously dominant in the sense of Section 5.1, since with a single contingency the worst case and the best case of each strategy coincide. The failure comes instead from the neighbor relation, under which the best strategy is two steps from strategy 0. With a single contingency, local improvement solvability is the condition that every local maximum of the payoff under the neighbor relation is a global maximum. The neighbor relation thus captures a source of simplicity beyond contingent reasoning, one that concerns the search through the strategy set. With three strategies the search is short, since the participant can compare every pair. Section 6.3 discusses Sudoku, in which the strategies are the many fillings of a grid and the same failure seems to make the puzzle difficult.

Local improvement solvability depends on the neighbor relation, and the dependence is monotone in the order of Definition 2.

Proposition 1: (a) *If Γ is local improvement solvable for θ_i under N_i and N_i is finer than N'_i , it is local improvement solvable for θ_i under N'_i .*

(b) *Γ is local improvement solvable for θ_i under the complete neighbor relation if and only if θ_i has a dominant strategy.*

The proof is in Appendix A. Since every neighbor relation is finer than the complete one, the proposition orders the family of conditions by the neighbor relation, with strategy-proofness

as its coarsest member. Strategy-proofness uses no structure on the strategy set, and under the complete neighbor relation the participant compares the candidate strategy against every other strategy. A finer neighbor relation uses more of the structure and gives a stronger condition. The strengthening can be strict, as the mechanism in Theorem 1(b) shows.

Which neighbor relation should the analysis use? I treat the neighbor relation as part of the description of a mechanism, much as a type space is part of the description of an environment. In each application I take the relation that the presentation of the strategy set defines, which I call its natural neighbor relation. Under the natural neighbor relation, the neighbors of a strategy are the strategies one change away in its presentation, and for a bid they are the one-step neighbors.¹⁰ Choosing the natural neighbor relation makes local improvement solvability as demanding as the presentation allows and the comparison between two neighbors as small as it can be.

Appendix C.1 explains the separate roles of the two requirements in Definitions 5 and 8, that a local improvement be strictly better in some contingency and that all locally unimprovable strategies be payoff-equivalent.

4 Classical mechanisms are local improvement solvable

Are the classical strategy-proof mechanisms local improvement solvable? This section shows that five of them are, under their natural neighbor relations: the sealed-bid second-price auction, the multi-unit Vickrey auction, serial dictatorship, top trading cycles, and deferred acceptance. For the three matching mechanisms, the result holds in a rich environment, defined in Section 4.3. In all five, the comparison between two neighboring strategies reduces to one case whose sign the participant determines from her own preference. The mechanisms are direct, apart from the grid of bids in the two auctions, and I call the recommended strategies truthful. All five are static mechanisms, in which each participant submits her strategy in one move and the outcome is computed from the strategies. Their dynamic counterparts are the subject of Section 7.

4.1 The sealed-bid second-price auction

The sealed-bid second-price auction is defined as follows. There are $n \geq 2$ bidders and one object. Each bidder submits a bid on the grid $\{0, \dots, K\}$ with $K \geq 2$, simultaneously and without observing the others' bids. The object goes to the bidder with the highest bid, who

¹⁰The relations for the other mechanisms are defined where the mechanisms are introduced, in Sections 4 to 7.

pays the highest bid among the other bidders. If two or more bidders tie for the highest bid, the object is not sold and no one pays. A bidder who does not win pays nothing.

Fix one bidder, the agent. Her type is her value θ for the object, with $0 \leq \theta < K$, which may or may not be a grid point. Her utility is her value minus her payment if she wins and zero otherwise. The contingency is the profile of the $n - 1$ rival bids. The agent's payoff depends on the rivals only through the highest rival bid p , and

$$u(b, p; \theta) = (\theta - p) \mathbf{1}[b > p],$$

where $\mathbf{1}[\cdot]$ is the indicator function. Let $T(\theta)$ be the set of grid points b with $\theta \leq b \leq \theta + 1$, the truthful bids. If θ is not a grid point, $T(\theta)$ contains one bid, the least grid point above θ . If θ is a grid point, $T(\theta) = \{\theta, \theta + 1\}$, and these two bids give the same utility in every contingency, since they give her different outcomes only when $p = \theta$. In that contingency the bid θ ties and loses, and the bid $\theta + 1$ wins at price θ , both with payoff 0.¹¹ The neighbor relation is one-step neighbors, so the neighbors of b are $b - 1$ and $b + 1$ when these are on the grid.

Proposition 2: *In the sealed-bid second-price auction, $u(b+1, p; \theta) - u(b, p; \theta) = (\theta - b) \mathbf{1}[p = b]$ for adjacent bids. Under one-step neighbors, the locally unimprovable strategies are the bids in $T(\theta)$, and the auction is local improvement solvable.*

The proposition follows directly from the payoff. Two neighboring bids give the agent different outcomes in exactly one case, the contingencies in which the highest rival bid equals the lower of the two bids. In that case the higher bid wins at that price and the lower bid loses. The difference is $\theta - b$, which the agent computes from her own value. It is positive for $b < \theta$, negative for $b > \theta$, and zero for $b = \theta$, the last case arising only when θ is a grid point.

Figure 1 shows the local improvements. An arrow from b to b' means that b' locally improves b , $\Delta_b = u(b + 1, p; \theta) - u(b, p; \theta)$ is the payoff difference between adjacent bids, and a^* is the least grid point at or above θ . The figure is drawn for a value that is not a grid point. Every bid below a^* is improved by the bid one step above it, and every bid above a^* by the bid one step below it. Thus a^* is the only locally unimprovable strategy and, by Theorem 1, the dominant bid. When θ is a grid point the arrow from $a^* + 1$ to a^* is absent, and a^* and $a^* + 1$ are two payoff-equivalent locally unimprovable strategies. In both cases the locally unimprovable strategies are the bids in $T(\theta)$, and every run of the algorithm ends at one of them.

The argument does not depend on K or on the number n of bidders, and the agent never compares a truthful bid against a bid other than its neighbors. By contrast, the sealed-bid

¹¹This setup includes the design of Li (2017) and Breitmoser and Schweighofer-Kodritsch (2022), in which values and bids are on the same grid and a tie for the highest bid means that no bidder wins.




Figure 1: Local improvements in the sealed-bid second-price auction, Proposition 2.

second-price auction is not OSP, and a direct comparison of a truthful bid with a bid farther away requires her to check every highest rival bid between the two.

4.2 Multi-unit demand and the Vickrey auction

The multi-unit Vickrey auction is defined as follows. There are Q identical units for sale and n bidders. Each bidder submits a weakly decreasing vector of marginal bids on the grid $\{0, \dots, K\}$, one bid per unit she may receive, simultaneously and without observing the others' bids. I assume that any $n - 1$ of the bidders together submit at least Q marginal bids. A bidder receives her ℓ -th unit exactly when her ℓ -th marginal bid exceeds the $(Q - \ell + 1)$ -th highest marginal bid among the other bidders. Without ties the units go to the Q highest marginal bids, and with ties some units can remain unsold, as the object does in the single-unit auction. A bidder who receives x units pays, for her ℓ -th unit, $\ell = 1, \dots, x$, the opportunity cost of that unit, the rival marginal bid that it displaces. The displaced bid is the $(Q - \ell + 1)$ -th highest marginal bid among the other bidders. This payment rule is the Vickrey, or VCG, pricing rule for identical units, and it reduces to the second-price rule when $Q = 1$.

Fix the agent. She demands up to $L \leq Q$ units. Her type is the vector $\theta = (\theta^1, \dots, \theta^L)$ of marginal values, with $K > \theta^1 > \dots > \theta^L \geq 0$, on or off the grid. Her utility is the sum of the marginal values of the units she receives minus her payments. Her strategy is a weakly decreasing vector $a = (a^1, \dots, a^L)$ of marginal bids on the grid. For each unit ℓ let $T(\theta^\ell)$ be the set of grid points b with $\theta^\ell \leq b \leq \theta^\ell + 1$, as in the single-unit auction, and call a strategy truthful if $a^\ell \in T(\theta^\ell)$ for every ℓ . Since the marginal values are strictly decreasing, every truthful strategy is weakly decreasing. The contingency is the profile of the rivals' strategies. Under the pricing rule just stated, the agent's payoff depends on the rivals only through the vector of residual prices $p = (p_1, \dots, p_L)$, where p_ℓ is the $(Q - \ell + 1)$ -th highest rival bid, so that $p_1 \leq \dots \leq p_L$. The residual prices are functions of the rivals' bids alone, and by the assumption above each can take every grid value. The agent wins unit ℓ if and only if $a^\ell > p_\ell$, and pays p_ℓ for it. Since a is weakly decreasing and p is weakly increasing, the units she wins

are units 1 through x for some x . Her payoff is

$$u(a, p; \theta) = \sum_{\ell=1}^L (\theta^\ell - p_\ell) \mathbf{1}[a^\ell > p_\ell], \quad (2)$$

a sum of L single-unit second-price payoffs. The neighbors of a strategy are its one-step neighbors on vectors, the weakly decreasing vectors that differ from it in one unit's bid by one grid step.

Proposition 3: *In the multi-unit Vickrey auction, raising the bid for unit ℓ from a^ℓ to $a^\ell + 1$ changes the payoff by $(\theta^\ell - a^\ell) \mathbf{1}[p_\ell = a^\ell]$. Under one-step neighbors on vectors, the locally unimprovable strategies are the truthful strategies, and the auction is local improvement solvable.*

The proof is in Appendix A. The multi-unit problem consists of L copies of the single-unit problem, connected only by the constraint that the strategy is weakly decreasing. The proof shows that from every strategy that is not truthful, the constraint leaves at least one local improvement toward a truthful strategy. As in the single-unit auction, the comparison between two neighboring strategies involves one case, the contingencies in which unit ℓ 's residual price equals the lower bid. The sign of the payoff difference in that case is determined by unit ℓ 's marginal value.¹²

4.3 Matching: local comparisons on rankings

A matching mechanism assigns each of n agents one of m objects. Object o has a capacity $q_o \geq 1$, the number of agents it can be assigned to, and $\sum_o q_o \geq n$, so that every agent is assigned an object. Each agent submits a strict ranking of all m objects. Each object o has a strict priority ranking $\succ_o$ over the agents. In the school choice interpretation the objects are schools with seats, and in the housing interpretation each object has capacity one and its owner has top priority at it. The three mechanisms studied are the following.

Serial dictatorship. The agents are ordered by a single priority order, the same for every object, which may be drawn at random. The first agent in the order receives the object she ranks highest, the second receives the object she ranks highest among those with remaining capacity, and so on until every agent is assigned.

¹²Identity (2) holds under Vickrey pricing. It does not hold under uniform pricing, where the price of every unit depends on the agent's own strategy, so the argument does not apply to uniform-price auctions. The Ausubel auction of Ausubel (2004) is a dynamic implementation of the Vickrey rule, and Section 7 discusses it.

Top trading cycles. In each round, every object with remaining capacity points to the remaining agent with the highest priority at it. Every remaining agent points to the object she ranks highest among those with remaining capacity. Since every remaining agent and every remaining object points at something, there is at least one cycle, and no two cycles intersect. Every agent in a cycle is assigned the object she points to and is removed, the capacity of that object falls by one, and the round ends. The mechanism stops when every agent is assigned. With capacities of one and owners at the top of the priorities this is the top trading cycles mechanism for housing markets of [Shapley and Scarf \(1974\)](#).

Deferred acceptance. In the first round every agent applies to the object she ranks highest. Each object tentatively holds the applicants with the highest priority at it, up to its capacity, and rejects the rest. In each later round every agent rejected in the previous round applies to the object she ranks highest among those that have not rejected her. Each object considers its new applicants together with the applicants it holds, again holds those with the highest priority up to its capacity, and rejects the rest. The mechanism stops in the first round with no rejection, and every agent is assigned the object that holds her. This is the agent-proposing, or student-proposing, deferred acceptance mechanism of [Gale and Shapley \(1962\)](#). Since every agent ranks every object and capacities suffice, every agent is assigned.

Fix an agent. Her type θ is her true strict ranking of the objects, and her strategy s_i is a strict ranking. The contingency s_{-i} is the profile of the others' strategies together with the priority data of the mechanism. The priority data are the priority order in serial dictatorship and the rankings $\succ_o$ in top trading cycles and deferred acceptance. Write $f(s_i, s_{-i})$ for the object she receives. The full strict domain is the type space in which every strict ranking is a possible type. The neighbors of a list are its one-exchange neighbors, the $m - 1$ lists obtained by exchanging two entries in adjacent positions.

Definition 9 (Adjacent-swap property): *The mechanism has the adjacent-swap property if, whenever s_i and s'_i differ by an adjacent exchange of o and o' with s_i listing o first, for every s_{-i} either $f(s_i, s_{-i}) = f(s'_i, s_{-i})$ or $f(s_i, s_{-i}) = o$ and $f(s'_i, s_{-i}) = o'$. In the second case the pair $\{o, o'\}$ is decisive at s_{-i} in both s_i and s'_i .*

Definition 10 (Rich environment): *The environment is rich if every adjacent pair of every list is decisive at some contingency.*

Under the adjacent-swap property, two strategies that differ by an adjacent exchange of o and o' lead to different outcomes only in contingencies where the pair is decisive. There the strategy listing o first receives o and the other receives o' . The comparison between the two strategies therefore involves one case. The sign of the payoff difference in that case is

determined by whether θ ranks o above o' . The exchange toward θ is a local improvement if the pair is decisive in at least one contingency, which is what richness provides.

Proposition 4: *Let strategies be strict rankings with one-exchange neighbors.*

(a) *If f has the adjacent-swap property and the environment is rich, then θ is the only locally unimprovable strategy, and f is local improvement solvable under one-exchange neighbors.*

(b) *Every deterministic mechanism on the full strict domain that is strategy-proof under truthful reporting has the adjacent-swap property.*

The proof is in Appendix A. For part (a), a list other than θ has an adjacent pair ordered differently from θ . Exchanging that pair is a local improvement by the adjacent-swap property and richness, and each exchange reduces the number of pairs ordered differently from θ by one, so every run ends at θ . For part (b), strategy-proofness for the two types whose true rankings are the two lists forces any change in the outcome to be the exchange of the two objects.

Corollary 1: *Serial dictatorship, top trading cycles, and student-proposing deferred acceptance have the adjacent-swap property, and in a rich environment they are local improvement solvable under one-exchange neighbors.*

The corollary follows from Proposition 4. The three mechanisms are deterministic, and on the full strict domain they are strategy-proof under truthful reporting for the agents. For deferred acceptance, strategy-proofness is shown by Roth (1982a) and Dubins and Freedman (1981), and for top trading cycles by Roth (1982b) in housing markets and by Abdulkadiroğlu and Sönmez (2003) with capacities and priorities. Part (b) therefore gives the adjacent-swap property. For serial dictatorship the property also follows from the rules directly. The earlier dictators do not read the agent's list, so the exchange of o and o' changes her outcome only when both are still available and she would have taken o . In a rich environment, part (a) then applies.

For example, the environment is rich when the agent's priority is part of the contingency, every profile of priorities and others' strategies is possible, and there are enough other agents. Appendix A.5 verifies richness under these conditions for each of the three mechanisms.

Richness depends on what the agent knows, and it fails when she knows that some adjacent pair on her list cannot affect her outcome. A list that is not payoff-equivalent to θ can then have no local improvement, as an example in Appendix A.5 shows.

In the random serial dictatorship treatments of Li (2017), every position remains possible whatever a subject's priority score, so the environment is rich and Corollary 1 applies. Appendix B.5 shows that the exchange of her two highest entries can be decisive whenever she

does not move last, and this exchange is the most common error that Li (2017) reports under list submission.

By Proposition 4(b), the one-case comparison is available in every deterministic ranking mechanism on the full strict domain that is strategy-proof under truthful reporting. What differs across mechanisms is how the participant can learn that an adjacent exchange changes nothing but the two exchanged objects. In serial dictatorship this follows from the rules. In top trading cycles and deferred acceptance it is a consequence of strategy-proofness, and Section 7.3 and Appendix B.4 argue that dynamic mechanisms display it.

5 Extensive-form mechanisms and games: locality along situations

The neighbor relations of Section 4 are defined on the strategies of static mechanisms. This section treats extensive-form mechanisms formally, and with them extensive-form games. What is the neighbor relation there, and what is local improvement solvability under it? Section 5.1 defines extensive-form mechanisms, one-situation neighbors, and four conditions, three existing and one introduced in this paper. Section 5.2 relates local improvement solvability under one-situation neighbors to strict action-proofness, and Section 5.3 relates the four conditions to one another.

5.1 Extensive-form mechanisms and simplicity concepts

Consider a finite extensive-form mechanism of perfect recall, a game tree with an outcome function, with nodes defined as action sequences as in Osborne and Rubinstein (1994). Let I be the set of players, including nature as the nonstrategic player 0, and let $I_+ = I \setminus \{0\}$ be the set of strategic players. I follow the terminology of Watson and Zhang (2026). A situation is a place in the game where a player must choose an action. I reserve the term information set for the information about the strategy profile that the player has at that place, expressed as a set of strategy profiles. Let H be the set of situations, let H_i be the set of situations of player $i \in I$, and for $h \in H$ let $A(h)$ be the set of actions available at h , with $|A(h)| \geq 2$. A strategy profile is $s = (a_h)_{h \in H}$ with $a_h \in A(h)$, and the set of strategy profiles is $S = \times_{h \in H} A(h)$. For $L \subset H$, the profile of actions at the situations in L is $s_L = (a_h)_{h \in L}$, with S_L the corresponding projection of S and $-L = H \setminus L$. Player i 's strategy is $s_i = s_{H_i}$, and $S_i = \times_{h \in H_i} A(h)$ is her strategy set. A strategy specifies an action at every situation in H_i , including situations that the strategy itself makes unreachable. For $h \in H_i$, s_h is the action that s_i specifies at h .

Nature's strategy s_0 is a component of the profile, and no probability distribution over nature's strategy is specified. A contingency for player $i \in I_+$ is a profile $s_{-i} \in S_{-i}$ of the other players' strategies, nature's included. Player i has a type $\theta_i \in \Theta_i$ and a payoff function $u_i : S \times \Theta_i \rightarrow \mathbb{R}$, written $u_i(s_i, s_{-i}; \theta_i)$, which is her utility from the outcome at the terminal node that the profile reaches.

The extensive-form mechanism is a mechanism in the sense of Section 3.1, with the strategy sets S_i and the outcome function given by the terminal nodes. Its reduced normal-form mechanism is the mechanism obtained by identifying two strategies of a player whenever they reach the same terminal node in every contingency. A static mechanism is an extensive-form mechanism in which every strategic player has exactly one situation, and a dynamic mechanism is one in which some strategic player has more than one. When the payoffs are fixed rather than indexed by types, the same object is called a game. An extensive-form game is therefore the special case of an extensive-form mechanism in which every player has a single type.

Conversely, fixing a type θ_i of player i turns an extensive-form mechanism into an extensive-form game for her, with payoffs $u_i(\cdot; \theta_i)$ and with the other players and nature as the source of her contingencies. Each condition in this section, and local improvement solvability, is stated for a single type and quantifies only over that type's strategies and the contingencies she faces. A mechanism therefore has each property exactly when every game that a type faces has it. The definitions and results below apply without change to extensive-form games, and the examples that separate the conditions are games of this kind.

For $h \in H$, let $S(h)$ be the set of strategy profiles that reach h , that is, the profiles whose induced action sequence has a node in h as a truncation. $S(h)$ is the information set at h . For $i \in I_+$ and $h \in H_i$, perfect recall implies that $S(h)$ is a restriction for player i , $S(h) = S(h)_i \times S(h)_{-i}$. A strategy s_i reaches h if $s_i \in S(h)_i$, and a contingency s_{-i} reaches h if $s_{-i} \in S(h)_{-i}$. A situation $h \in H_i$ precedes $h' \in H_i$, and h' follows h , if every node in h' has a node in h as a strict prefix. Perfect recall implies that for every $h' \in H_i$ that precedes h the projection $S(h)_{h'}$ is a single action, the action at h' that leads toward h . Perfect recall also implies that for every other $h' \in H_i$ the projection $S(h)_{h'}$ is $A(h')$. Hence $S(h)_i$ is the set of strategies that specify, at every situation of player i preceding h , the action that leads toward h . Two strategies that both reach h therefore agree at every situation preceding h . For a strategy $r \in S_i$, a situation $h \in H_i$, and an action $a \in A(h)$, write $r[h \leftarrow a]$ for the strategy $(a, r_{H_i \setminus \{h\}})$ that agrees with r at every situation other than h and specifies a at h . A continuation of an action a at h is a strategy $r \in S(h)_i$ with $r_h = a$. A situation h' follows h through a if h' follows h and a is the action at h that leads toward h' , so that $S(h')_h = \{a\}$. The actions that r specifies at the situations that follow h through a are its continuation after a .

The information in the extensive-form structure already restricts the comparisons that a player must make. At a situation h she knows that play has reached h , so the strategies she compares there are those in $S(h)_i$. Every strategy that reaches h is then in effect a neighbor of every other strategy that reaches h . The contingencies are reduced in the same way. Suppose two strategies differ only at h , or only at h and at situations following h . In every contingency that does not reach h they lead to the same terminal node, since play never passes through a situation at which they differ. A comparison between the two therefore involves only the contingencies in $S(h)_{-i}$. The extensive-form mechanism shows the player both restrictions. Its reduced normal-form mechanism presents the same strategies as wholes and shows neither. When $S(h)_{-i}$ is smaller than S_{-i} , the dominant strategy can therefore be easier to find in an extensive-form mechanism than in its reduced normal-form mechanism. The conditions defined below all use these restrictions.

A finer relation can also be used. The one-situation neighbors of s_i are the strategies that differ from s_i at exactly one situation,

$$N_i(s_i) = \{s'_i \in S_i : |\{h \in H_i : s_h \neq s'_h\}| = 1\}.$$

Under this relation, a step of the local improvement algorithm changes the action at one situation and holds the rest of the strategy fixed. The one-situation neighbors of s_i are thus its one-shot deviations $s_i[h \leftarrow a]$ with $h \in H_i$ and $a \in A(h) \setminus \{s_h\}$, the strategies against which the one-shot deviation principle checks s_i .¹³ Section 5.2 shows what local improvement solvability then requires, and Section 5.3 ends by comparing it with the principle. One-situation neighbors differ in a single action, the smallest change of a strategy that the extensive-form structure defines.¹⁴

The following definitions state three existing concepts, namely dominance, the obvious dominance of Li (2017), and the strong obvious dominance of Pycia and Troyan (2023). The definitions also state a fourth condition, action dominance, which I introduce in this paper. Fix a player $i \in I_+$, a type θ_i , and a strategy s_i^* , called the recommended strategy. Under a recommendation σ in the sense of Section 3.1, $s_i^* = \sigma_i(\theta_i)$. A strategy s'_i deviates from s_i^* at a situation $h \in H_i$ if both strategies reach h , that is, $s_i^*, s'_i \in S(h)_i$, and $s'_h \neq s_h^*$. Recall that two strategies that reach h agree at every situation preceding h . Hence the situations at which s'_i deviates from s_i^* are those at which the two strategies first differ along some path of play, which Li (2017) calls their earliest points of departure. Each condition below is stated at a

¹³The principle is presented, for instance, in Osborne and Rubinstein (1994), and Hendon, Jacobsen and Sloth (1996) extend it from subgame perfection to sequential rationality.

¹⁴A mechanism whose strategy set has more structure can support a relation with an even simpler path to the dominant strategy, as the ascending clock auction of Section 7.1 shows.

situation h that s_i^* reaches, and the contingencies in each comparison are those that reach h , the elements of $S(h)_{-i}$. Each condition depends on h through $S(h)_{-i}$ and through the set $S(h)_i$ of the player's own strategies that reach h , which in action dominance and strong obvious dominance supplies both sides of the comparison.

Definition 11 (Dominance at a situation): *The strategy s_i^* is dominant for θ_i if (1) holds. Equivalently, at every situation h that s_i^* reaches and for every s'_i that deviates from s_i^* at h ,*

$$u_i(s_i^*, s_{-i}; \theta_i) \geq u_i(s'_i, s_{-i}; \theta_i) \quad \text{for every } s_{-i} \in S(h)_{-i}. \quad (\text{D})$$

Definition 12 (Obvious dominance): *The strategy s_i^* is obviously dominant for θ_i if, at every situation h that s_i^* reaches and for every s'_i that deviates from s_i^* at h ,*

$$\min_{s_{-i} \in S(h)_{-i}} u_i(s_i^*, s_{-i}; \theta_i) \geq \max_{s'_{-i} \in S(h)_{-i}} u_i(s'_i, s'_{-i}; \theta_i). \quad (\text{OD})$$

Definition 13 (Action dominance): *The strategy s_i^* is action dominant for θ_i if, at every situation h that s_i^* reaches,*

$$u_i(r[h \leftarrow s_h^*], s_{-i}; \theta_i) \geq u_i(r[h \leftarrow a'], s_{-i}; \theta_i) \quad (\text{AD})$$

for every $r \in S(h)_i$, $a' \in A(h)$, and $s_{-i} \in S(h)_{-i}$.

It is strictly action dominant if, in addition, for every such h , r , and $a' \neq s_h^*$ the inequality is strict for some $s_{-i} \in S(h)_{-i}$.

Definition 14 (Strong obvious dominance): *The strategy s_i^* is strongly obviously dominant for θ_i if, at every situation h that s_i^* reaches and for every $a' \in A(h)$ with $a' \neq s_h^*$,*

$$\min_{\substack{r \in S(h)_i: r_h = s_h^* \\ s_{-i} \in S(h)_{-i}}} u_i(r, s_{-i}; \theta_i) \geq \max_{\substack{r' \in S(h)_i: r'_h = a' \\ s'_{-i} \in S(h)_{-i}}} u_i(r', s'_{-i}; \theta_i). \quad (\text{SOD})$$

The mechanism is strategy-proof (SP), obviously strategy-proof (OSP), or action-proof (AP) if every type of every strategic player has a dominant, an obviously dominant, or an action dominant strategy, respectively. It is strictly action-proof or strongly obviously strategy-proof

(SOSP) if every type of every strategic player has a strictly action dominant or a strongly obviously dominant strategy, respectively.¹⁵

Condition (D) is dominance stated situation by situation. Condition (OD) is the definition of Li (2017), and condition (SOD) is the definition of Pycia and Troyan (2023). Compare the two definitions. In (OD) the deviation is a strategy, and the comparison holds the recommended strategy fixed after h . In (SOD) the deviation is an action, and the comparison ranges over everything the player might do after h on either side. These three conditions compare payoffs across the contingencies in $S(h)_{-i}$, either one contingency at a time or by the worst case against the best case, and none uses a neighbor relation on the strategy set. Conditions (OD) and (SOD) make the second comparison and are worst-best comparison concepts in the sense of the introduction.

Action dominance takes the deviation of (SOD), an action, and the comparison of (D), one contingency at a time.¹⁶ The player's strategy at her other situations is arbitrary but the same on both sides of the comparison. Action dominance also has a precedent in the agent normal form of Selten (1975), in which each situation of a player is treated as a separate agent who chooses the action there. Action dominance at h is weak dominance of the recommended action for the agent at h , with the player's own actions at her other situations among that agent's contingencies. This weak dominance is required at every situation that s_i^* reaches. Unlike conditions (D), (OD), and (SOD), action dominance compares two strategies that differ at one situation only. A change of the action at one situation is a move to a one-situation neighbor, a one-shot deviation. Theorem 2 shows that, in extensive-form mechanisms without duplicate dominant strategies, local improvement solvability under one-situation neighbors is equivalent to strict action-proofness.

Each of the four conditions is stated situation by situation, and Section 5.3 and Appendix C.2 use the version at a single situation. Fix a situation h that s_i^* reaches and an action $a' \in A(h)$ with $a' \neq s_h^*$. The strategy s_i^* satisfies dominance at h against a' if (D) holds for every $s'_i \in S(h)_i$ with $s'_h = a'$. It satisfies obvious dominance at h against a' if (OD) holds for every such s'_i . It satisfies action dominance at h against a' if (AD) holds for every $r \in S(h)_i$ and every $s_{-i} \in S(h)_{-i}$. It satisfies strong obvious dominance at h against a' if (SOD) holds. Appendix A.6 proves the equivalence of (1) and (D) in Definition 11.

The last concept is the complexity measure of Nagel and Saitto (forthcoming). It first merges consecutive situations of the player between which no payoff-relevant information

¹⁵For two such properties, $X \Rightarrow Y$ means that every mechanism with X has Y , $X \not\Rightarrow Y$ that some mechanism with X lacks Y , and $X + Y$ is the conjunction. The same notation is used for the conditions (D), (OD), (AD), and (SOD) on a strategy.

¹⁶One precedent is the weak dominance between actions of Catonini and Xue (2026), which Section 5.3 compares with action dominance.




Figure 2: A game with duplicate dominant strategies, Theorem 2.

arrives. It then removes, one at a time, actions whose best payoff at a situation, over everything that can follow, is no larger than the worst payoff of another surviving action there. Finally, it counts the actions that survive, and the complexity is the largest such count over the players, their types, and their situations.¹⁷ The results in Section 6 use only two properties of the measure. The removal step compares the maximum of one set of payoffs against the minimum of another. Two consecutive situations of the player are merged unless the information between them excludes some payoff possible under her dominant continuation.

5.2 One-situation neighbors

A change of action at a situation that the current strategy does not reach changes no outcome, so it is not a local improvement. A change of action at a situation h that the current strategy reaches replaces the strategy r by $r[h \leftarrow a]$. The comparison between the two strategies is the comparison in (AD) with r fixed, together with the requirement that the change be strictly better in some contingency. Strict action dominance, Definition 13, is exactly (AD) together with that requirement.

Two dominant strategies of a type are duplicates if they specify different actions at some situation that both of them reach. Two dominant strategies are payoff-equivalent, so duplicates exist only when two different actions at some situation lead to continuations that are equally good in every contingency, as in the game of Figure 2. In a static mechanism every strategy reaches the player's single situation, so any two different dominant strategies are duplicates. The truthful bids θ and $\theta + 1$ of the sealed-bid second-price auction with a value on the bid grid are an example.

Theorem 2: *The type θ_i has a strictly action dominant strategy if and only if Γ is local improvement solvable for θ_i under one-situation neighbors and θ_i has no duplicate dominant strategies.*

¹⁷Appendix A.2 gives a simplified version of the construction in the notation of this paper and defines the terms used in Section 6, such as the evident reduction and the essential frame.

The proof is in Appendix A. If s_i^* is strictly action dominant, any strategy that differs from it at a situation both reach can be improved there. Every locally unimprovable strategy is therefore payoff-equivalent to s_i^* , and a similar argument rules out duplicates. Conversely, let s_i^* be dominant and suppose strict action dominance fails for it at a situation h . A run of the algorithm from the failing deviation at h , with the actions after s_h^* held fixed, ends at a locally unimprovable strategy that does not play s_h^* at h . Under local improvement solvability this strategy is dominant and differs from s_i^* at a situation both reach, so the two are duplicates.

Consider the dynamic serial dictatorship of Section 7.3, in which the agents are called one at a time in the priority order and each takes one of the objects that remain. Each agent's choice from the offered set determines her outcome, so choosing her most preferred object in that set is strictly action dominant. By the theorem, the dynamic serial dictatorship is local improvement solvable under one-situation neighbors. By Corollary 1, its direct mechanism, serial dictatorship with list submission, is local improvement solvable under one-exchange neighbors in a rich environment.

Local improvement solvability alone does not rule out duplicates. In the game of Figure 2, the agent chooses a^* or a' , and the two choices lead to identical subgames. After a^* she chooses x , paying 10, or y , paying 0, and after a' she chooses x' , paying 10, or y' , paying 0. Here and in Section 5.3, a strategy is written as the list of the agent's actions, from the top of the tree down and from left to right. A dot marks a situation whose action does not matter. Every strategy with payoff 0 is improved by correcting its action at the second situation it reaches, so every locally unimprovable strategy pays 10 and the game is local improvement solvable under one-situation neighbors. The strategies $(a^*, x, \cdot)$ and $(a', \cdot, x')$ are both dominant and differ at the first situation, which both reach, so they are duplicates. The game is not action-proof. If a^* is recommended, the strategy $r = (\cdot, y, x')$ pays 0 with a^* and 10 with a' , and the case with a' recommended is symmetric.

5.3 Four conditions on a step

The four conditions of Definitions 11 to 14 differ in two respects. The first is whether the player's own actions after the recommended action are held equal to those of s_i^* or allowed to vary. The second is whether the contingencies are compared one at a time or by comparing the worst case of one side against the best case of the other. Table 1 arranges the four conditions. Its rows are the two ways of comparing contingencies and its columns are the two treatments of the player's own other actions.

Moving down a column replaces the comparison contingency by contingency with the comparison of the worst case against the best case. This replacement is the difference between




Figure 3: The two-move game, Proposition 5(b).

	own other actions equal to s_i^*	own other actions arbitrary
contingency by contingency	strategy-proof	action-proof
worst case against best case	OSP	SOSP

Table 1: Four conditions on a step of the local improvement algorithm at a situation.

strategy-proofness and OSP and between action-proofness and SOSP, and it removes the need to consider which contingency obtains. Moving across a row allows the player’s own actions at other situations to vary. This variation removes the need to consider what the player will do at later situations. Action-proofness is the condition obtained by moving across a row alone.

Most other worst-best comparison concepts lie between these four conditions. The partition-based concepts of [Zhang and Levin \(2018\)](#) and [Arribillaga, Massó and Neme \(2024\)](#) fill the left column of Table 1, between dominance and obvious dominance, and the concepts of [Pycia and Troyan \(2023\)](#) fill the bottom row. The conditions of [Catonini and Xue \(2026\)](#) compare two actions and treat the player’s own continuation in a third way. Appendix C.2 relates all of them to the table.

Comparing actions is not what places a condition in the right column, since the conditions of [Catonini and Xue \(2026\)](#) compare actions as well. What places it there is the requirement that the comparison hold for every continuation after the recommended action. Action dominance is the entry of that column in which the contingencies are compared one at a time. As a property of mechanisms it does not appear in the literature as far as I know, although each neighboring condition in the table does. The weak dominance between actions of [Catonini and Xue \(2026\)](#) is a closely related precedent. It asks only that some continuation after the recommended action match each continuation after a' . Weak dominance between actions, required at every situation that a strategy reaches, is equivalent to dominance. Hence the requirement for every continuation after the recommended action is exactly what separates action-proofness from strategy-proofness. I know of no concept between strategy-proofness and action-proofness, or between action-proofness and SOSP.

Proposition 5: (a) $(SOD) \Rightarrow (OD) \Rightarrow (D)$ and $(SOD) \Rightarrow (AD) \Rightarrow (D)$. Hence $SOSP \Rightarrow OSP \Rightarrow SP$ and $SOSP \Rightarrow AP \Rightarrow SP$.

(b) No implication in (a) reverses, $OSP \not\Rightarrow AP$, and $AP \not\Rightarrow OSP$.

(c) $OSP + AP \not\Rightarrow SOSP$.

Parts (a) and (c) together say that obvious strategy-proofness and action-proofness are both necessary for strong obvious strategy-proofness, and that the two together are not sufficient for it. The proof is in Appendix A. All four implications in part (a) are immediate from the definitions, situation by situation.

Parts (b) and (c) are proved by four examples. The two that show that obvious strategy-proofness and action-proofness are independent are given here, and the other two are in Appendix A. In the game of Figure 3 no other player moves, so there is a single contingency, and strategies are written as in Section 5.2.

The game of Figure 3 is OSP and not AP. The recommended strategy $(a^*, x, \cdot)$ pays 10, the best payoff in the game, so it is obviously dominant. It is not action dominant, since for the strategy $r = (\cdot, y, y')$ the first action a^* pays 0 and a' pays 5. The recommended action at the first situation is best only if the agent goes on to play x . Obvious dominance holds this later action fixed, and action dominance does not.

The sealed-bid second-price auction is AP and not OSP. In a static mechanism each player has one situation, at which action dominance and dominance are the same condition, so the auction is action-proof, and Li (2017) shows that it is not OSP.

Theorem 2 places local improvement solvability under one-situation neighbors next to action-proofness in Table 1. Absent duplicate dominant strategies, the two differ only by strictness, so local improvement solvability under one-situation neighbors relates to obvious strategy-proofness as action-proofness does. The sealed-bid second-price auction with a value off the grid is strictly action-proof and not OSP, while the two-move game of Figure 3 is OSP and not action-proof. Strong obvious strategy-proofness implies action-proofness but not local improvement solvability under one-situation neighbors. In dynamic Sudoku, Section 6.3, the recommended digit at a cell is action dominant without being strictly so, since an error at a later cell can leave the payoff at 0 whichever digit is entered. This lack of strictness is the only way a strongly obviously strategy-proof mechanism can fail to be local improvement solvable under one-situation neighbors. The two-move game of Figure 3 shows the other way a second locally unimprovable strategy arises. The strategy (a', y, y') pays 5 and its three one-situation neighbors pay 0, 5, and 0, so it can be improved only by changing a' to a^* and y to x at once.

Theorem 2 can be read as a one-shot deviation principle for dominance. The classical principle fixes the strategies of the others and checks the one-shot deviations from the candidate

strategy at every situation, as if play had reached it. In a finite game, a strategy that passes every such check is optimal. Action dominance requires that, at each situation that s_i^* reaches, no one-shot deviation from the recommended action pay more, whatever the player does at her other situations and whatever the others do. When the type has a strictly action dominant strategy, every strategy with no improving one-shot deviation is dominant, by the theorem and Theorem 1(a). Absent duplicate dominant strategies, the converse holds as well. The requirement on the player's other actions matters because the local improvement algorithm may start from a strategy that errs at several situations, such as (a', y, y') .

6 Three games and the worst-best comparison concepts

This section examines the worst-best comparison concepts through three games. The pick-a-number game, which a short argument solves, is not OSP and is maximally complex in the measure of Nagel and Saitto (forthcoming). Chess, when White has a winning strategy, is OSP for White, and Sudoku with a unique solution is SOSP, although both seem difficult for a participant. Local improvement solvability under the natural neighbor relation of each game accounts for part of these discrepancies, since it holds in the easy game and fails in the two difficult ones.

6.1 The pick-a-number game

Two players each pick a number from $A = \{0, \dots, K\}$, and player i receives $u_i(a_i, a_j) = a_i + \lambda a_j$, where a_j is the opponent's number and $\lambda > K$. There is no private information, so the type is degenerate and omitted. The one-step neighbors of a number are the numbers one step above and below it. The classifications below hold the payoffs and the set A fixed.

Proposition 6: *In the pick-a-number game with $K \geq 1$ and $\lambda > K$:*

(a) $u_i(a + 1, a_j) - u_i(a, a_j) = 1$ for every $a < K$ and every a_j , and the game is local improvement solvable under one-step neighbors.

(b) The static game is not OSP, and its complexity in the measure of Nagel and Saitto (forthcoming) is $K + 1$.

(c) Let an SP extensive-form game implement the outcomes of the static game through maps ϕ , in the sense of Definition 15. For numbers $a \neq a'$ and b , let h be the first situation of player i on the play of $(\phi_i(a), \phi_j(b))$ at which $\phi_i(a)$ and $\phi_i(a')$ differ. Then every $\phi_j(b')$ reaches h . No such game is OSP. If each $\phi_i(a)$ gives i the number a whatever j does, the complexity is at least $K + 1$, and every strategy of j reaches every situation of i that two such strategies reach.

Parts (a) and (b) are immediate, and part (c) is proved in Appendix A. In part (a), the term λa_j cancels from the difference between two neighboring numbers, so the difference is 1 in every contingency. Hence $a + 1$ locally improves a for every $a < K$, and K is the only locally unimprovable strategy. The participant does not need to know which contingency obtains. Part (a) is a statement about payoff differences. The worst-best comparison concepts compare payoff levels. In part (b) the minimum payoff from K is K , at $a_j = 0$, and the maximum payoff from any other number a is $a + \lambda K > K$, at $a_j = K$, so (OD) fails. Since $a + \lambda K$ exceeds every number in A , no number is evidently never better than another, and the evident reduction is A .

For part (c), an extensive-form game can reveal a player's earlier actions to the other player, and the other player can condition on them. Since $\lambda > K$, such conditioning lets a player punish the opponent's recommended strategy. So wherever two of a player's strategies that represent numbers first differ on a play against such a strategy of the opponent, strategy-proofness requires every such strategy of the opponent to reach that situation. Along these strategies every strategy-proof implementation thus works like the static game, and (OD) fails where those for K and a smaller number first differ. The measure of Nagel and Saitto (forthcoming) also registers information about moves that represent no number, such as a veto. Suppose instead that each of a player's strategies that represents a number gives her that number whatever the opponent does. She then cannot tell two strategies of the opponent apart at a situation that two of these strategies reach, for the same reason. These situations therefore merge into one, so dynamic implementation cannot lower the complexity, although the assumption itself restricts no information.

The parameter λ shows how the classification by the worst-best comparison concepts depends on the weight of the opponent's number. For $\lambda = 1$ there is a dynamic game that implements the outcomes of the static game. In it the players alternately raise their numbers by one step, observe each other, and cannot resume raising once they stop. That game is obviously strategy-proof, because a number already committed bounds the loss from any later action of the opponent. For $\lambda > K$ that bound never suffices, and no strategy-proof extensive-form game that implements these outcomes is OSP. A classification by obvious strategy-proofness that changes with the exchange rate λ , while the decision is the same, seems to measure something other than the difficulty of the decision.

6.2 Chess

Chess and Sudoku show the opposite error. Consider chess with payoff 1 for a win, $1/2$ for a draw, and 0 for a loss, and suppose, following Pycia and Troyan (2023), that White has a

winning strategy.¹⁸ Any winning strategy is then obviously dominant, since its worst outcome at every position it reaches is a win and no deviation does better. It is not strongly obviously dominant, since after the winning move some continuation of White loses while some deviation wins against an error of Black.¹⁹ Under one-situation neighbors, however, chess is not local improvement solvable for White. Call a move losing if Black can prevent a win after it. Fix a winning strategy, a position p that it reaches, its move w there, and a losing move m at p after which White can still win against some error of Black. Take a strategy that makes m at p and a losing move at a position q that the winning strategy reaches after w . Appendix A.10 specifies the winning strategy and the other moves so that no other change at one position improves the strategy. Correcting m alone is not a local improvement, because the corrected strategy reaches q . The corrected strategy fails to win against a strategy of Black that leads to p , errs after m so that the original strategy wins, and after w leads to q and prevents a win. Correcting the move at q alone changes nothing, since q is not reached. The strategy with the two losing moves is therefore locally unimprovable, and it is not dominant, since it does not win against a strategy of Black that leads to p and plays correctly after m .

6.3 Sudoku

Sudoku is a single-agent game in which, as in chess, a strategy with two mistakes can have no improving neighbor. Let C be a set of cells and D a set of digits, in the usual puzzle the 81 squares of a 9×9 grid and $\{1, \dots, 9\}$. A filling is a function $g : C \rightarrow D$. I restrict attention to puzzles with a unique solution, so there is exactly one valid filling g^* , and the payoff of a filling is $w(g) = \mathbf{1}[g = g^*]$. There are no other players and no moves of nature, so there is a single contingency. I consider two games. In static Sudoku the agent submits a filling in one move, so the game has one situation, at which the actions are the fillings. The one-cell neighbors of a filling are the fillings that differ from it at one cell. In dynamic Sudoku the cells are filled in a fixed order, one cell per move, with the digits D as the actions and no information revealed between moves other than the agent's own entries. By perfect recall, a situation is a cell together with the digits already entered. A strategy specifies a digit at every situation, and each strategy leads to one filling, so the two games have the same reduced normal-form game.

Proposition 7: (a) *Static and dynamic Sudoku are SOSP, and their complexity in the measure of Nagel and Saitto (forthcoming) is 1.*

(b) *Static Sudoku is not local improvement solvable under one-cell neighbors, and dynamic Sudoku is not local improvement solvable under one-situation neighbors.*

¹⁸Black need not have a dominant strategy, so the classifications of chess in this paper concern White.

¹⁹Pycia and Troyan (2023) use this observation to motivate their stronger concepts.

The proof is immediate from the fact that every wrong filling has payoff 0 and g^* has payoff 1. For (a), the minimum in (SOD) after the correct digit at any cell, or after g^* in static Sudoku, is at least the maximum after a wrong one, since that maximum is 0. In dynamic Sudoku no information about the contingency arrives between moves, so its essential frame is static Sudoku. Every $g \neq g^*$ is also evidently never better than g^* . For (b), a filling with two or more wrong cells and all of its one-cell neighbors pay 0, so it is locally unimprovable and not payoff-equivalent to g^* . In dynamic Sudoku, take a strategy that enters wrong digits at a cell c and at a later cell c' when all earlier entries are correct, and correct digits everywhere else. Its filling has one wrong cell, correcting c alone leads to the wrong digit at c' , and every other single change either leaves a wrong cell or is at a situation that is not reached. So it is locally unimprovable with payoff 0, while the recommended strategy pays 1.

Both classifications in (a) record that the participant needs no contingent reasoning about others and no reasoning about her own later moves. They do not record that the participant must find g^* , and that comparing neighboring fillings does not help her find it. The comparison does not help because a single change of a wrong filling leaves the payoff at 0 unless every other cell is already correct. The difficulty of finding g^* is the primitive complexity of [Babaioff, Dobzinski and Kupfer \(2023\)](#), the number of operations needed to find the best element of a menu, here the fillings.²⁰

The proposition has two consequences. The first concerns dynamic implementation. Dynamic Sudoku is SOSP, so the obviousness concepts classify it as the simplest kind of extensive-form game, and yet filling the cells one at a time does not seem to help. The comparison at a cell has the same payoff on both sides unless the rest of the strategy completes g^* , so it carries no information about the solution until then. Obvious strategy-proofness therefore does not seem sufficient for dynamic implementation to help.

The second consequence concerns what the two approaches measure. Static Sudoku is a single decision with one contingency and no later move. It is SOSP, the strongest of the obviousness concepts, because these concepts hold trivially for any single choice with one contingency. It nevertheless seems difficult, for a reason that has nothing to do with contingent reasoning or with planning one's own later moves, the only sources of difficulty the obviousness concepts measure. A concept that classifies static Sudoku as no more difficult than choosing between two objects does not appear to capture all of the difficulty of recognizing the dominant strategy. Local improvement solvability captures part of it. With a single contingency, the comparison of two neighboring strategies has no contingent reasoning to reduce. Local improvement solvability is then the condition that every local maximum of the payoff is a

²⁰The errors that [Enke and Shubatt \(2023\)](#) document in choices without strategic content seem to be of this kind.

global maximum, as Section 3.2 notes, and part (b) shows that neither version of Sudoku meets it. As in the example of Theorem 1(b), the neighbor relation thus captures a source of simplicity beyond contingent reasoning, one that concerns the search through the strategy set.

6.4 What the two approaches measure

	Worst-best comparison concepts	Local improvement solvable
Pick-a-number game, $\lambda > K$	not OSP, complexity $K + 1$	yes, one-step neighbors
Chess, White has a winning strategy	OSP for White	no, one-situation neighbors
Sudoku, static	SOSP, complexity 1	no, one-cell neighbors
Sudoku, dynamic	SOSP, complexity 1	no, one-situation neighbors

Table 2: Classification of the three games, Proposition 6, Section 6.2, and Proposition 7.

Table 2 collects the classifications of the three games. Both approaches describe ways to simplify the recognition of the dominant strategy, and they classify the three games differently. The worst-best comparison concepts, including the measure of Nagel and Saitto (forthcoming), classify the extensive-form game by the contingent reasoning it requires. The pick-a-number game requires it if payoff levels are compared, and chess and Sudoku do not require it. Local improvement solvability classifies a game by the payoff differences between neighboring strategies, and asks whether their signs lead to the dominant strategy. This question arises even with a single contingency, as in Sudoku, where no contingent reasoning is needed. In the pick-a-number game every difference is 1. In chess and Sudoku some strategies with two mistakes have no improving neighbor, although they are not dominant. The three games therefore suggest that the need for contingent reasoning is neither necessary nor sufficient for a mechanism to be difficult.

Local improvement solvability under the natural neighbor relation of each game thus accounts for part of what the worst-best comparison concepts miss. The explanation depends on that relation, since under the complete neighbor relation the pick-a-number game and Sudoku are local improvement solvable, being strategy-proof, and so is chess for White, who has a dominant strategy.

7 Salience of the local comparison

Section 4 gave the participant a path to the dominant strategy, in which each step compares two neighboring strategies and involves one case whose sign she determines from her own preference. The path helps her only if she notices it. A participant who is asked for her whole

strategy at once may never compare a strategy with its neighbors. When a mechanism instead asks her at each step to choose between two neighbors, the question itself leads her along the path. Two strategy-proof mechanisms can therefore implement the same outcomes, be local improvement solvable for the same reason, and still differ in how easily the local comparison is noticed. The prediction of this section is that the one in which the comparison is more salient produces more dominant-strategy play.

The prediction compares two mechanisms. I define the relation between them formally and leave the salience of the local comparison informal.

Definition 15 (Implementing the outcomes): *Let Γ and Γ' , with outcome functions f and f' , share players, type sets, and outcome space, and be strategy-proof under recommendations σ and σ' . The mechanism Γ implements the outcomes of Γ' if there are maps $\phi_i : S'_i \rightarrow S_i$, the identity for nature, such that*

$$f(s'_0, \phi_1(s'_1), \dots, \phi_n(s'_n)) = f'(s'_0, s'_1, \dots, s'_n) \quad \text{for every profile } s' \in S', \quad (3)$$

and $\sigma_i = \phi_i \circ \sigma'_i$ for every player.

In words, every strategy of Γ' is represented in Γ by a strategy that produces the same outcome against every profile of the others' strategies carried over by ϕ , and the recommended strategies correspond. The two mechanisms therefore produce the same outcome at every type profile and every move of nature, $f(s_0, \sigma(\theta)) = f'(s_0, \sigma'(\theta))$, and implement the same choice rule in the sense of [Li \(2017\)](#).

The relation need not be symmetric. The sealed-bid second-price auction and the ascending clock auction of [Section 7.1](#) implement each other's outcomes. The map ϕ sends a bid b to the strategy that exits at b , and a map in the other direction sends each strategy of the ascending clock auction to its exit price. A strategy that stays in at every price below K is sent to K . The dynamic serial dictatorship of [Section 7.3](#) implements the outcomes of serial dictatorship with list submission, with ϕ sending a list to the strategy that takes the top available object on the list. The converse fails, since the dynamic serial dictatorship has strategies whose outcomes no list reproduces. When Γ implements the outcomes of Γ' , the map ϕ_i carries the neighbor relation of Γ' to the image $\phi_i(S'_i)$. A run of the local improvement algorithm in Γ' is then a run in Γ along that image under one condition. Every contingency of Γ must give the participant, on that image, the payoffs of some contingency of Γ' . The condition holds in the applications below, since there the others' play that affects her payoff cannot depend on her own.

Between two mechanisms of which one implements the outcomes of the other, the local comparison is more salient in the one that leads the participant to compare a strategy with its

neighbors. A mechanism leads her to the comparison in one of two ways. It can ask her for a choice that contains a comparison between two neighboring strategies, as the ascending clock auction does at each price and the dynamic serial dictatorship at each menu. Or it can resolve her submitted strategy step by step in the order of the neighbor relation, as the passive clock of [Breitmoser and Schweighofer-Kodritsch \(2022\)](#) does. A mechanism that asks for a bid, a vector of marginal bids, or a ranking in one move and resolves it at once does neither, however its rules are described. I use the ordering by salience only where the design makes it clear. The ordering is clear among the sealed-bid second-price auction, the passive clock, and the dynamic clock, and between list submission and a mechanism that asks for one object at a time. Among multi-unit auctions it is clear for the static Vickrey auction, the Ausubel auction, and the multi-account clock of Appendix C.4. Among dynamic mechanisms, a question confined to two neighbors makes the comparison more salient than one that also offers strategies that are not neighbors.

Why is salience not formalized? It is a property of what the participant is shown and asked, not of the extensive-form mechanism. The worst-best comparison concepts take the extensive-form mechanism as their only input, and so does local improvement solvability once the neighbor relation is fixed. A formal measure of salience would need a model of what a participant attends to, and the evidence below does not require one.

The prediction does not deny that obvious strategy-proofness helps, since an OSP mechanism removes the need to reason about contingencies. The local comparison may be a second reason a participant finds the dominant strategy. I read the evidence below as evidence about the local comparison, with obvious strategy-proofness held fixed where the designs allow.

7.1 The sealed-bid second-price auction and the ascending clock auction

The ascending clock auction is defined as follows. There are $n \geq 2$ bidders and one object, and the agent's value is as in Section 4.1. A price starts at 0 and rises one grid step at a time through $\{0, \dots, K\}$. At each price q , every bidder who is still active chooses to stay in or to exit. A bidder who exits is inactive for the rest of the auction and observes nothing further, and an active bidder observes only the current price.²¹ If exactly one bidder stays in at a price $q < K$, she wins the object and pays q . If all active bidders exit at the same price, or if two or more bidders are still active when the price reaches K , the object is not sold and no one pays. A strategy of the agent specifies, at every price at which she could still be active, whether to stay in or to exit. The bidder who plans to exit at price b , staying in at every lower price,

²¹In the design of [Li \(2017\)](#) and in some treatments of [Breitmoser and Schweighofer-Kodritsch \(2022\)](#), active bidders also see when rivals exit, information that is payoff-irrelevant under private values and is discussed below.

obtains the object exactly when every rival exits before b . In that case she pays the price at which the last rival exits. She gets nothing when some rival exits at b or later. Her payoff is therefore that of the bid b in the sealed-bid second-price auction of Section 4.1, including the tie rule, so the two auctions implement each other's outcomes.

Strategy-proofness alone therefore does not distinguish the two auctions. Yet [Li \(2017\)](#) finds that the ascending clock auction produces more dominant-strategy play than the sealed-bid second-price auction and explains the difference by obvious strategy-proofness. The ascending clock auction is OSP. At each price, the worst outcome of the recommended action, staying in below the value and exiting at or above it, is at least as good as the best outcome of the other action. The sealed-bid second-price auction is not OSP. The ascending clock auction, however, also differs from the sealed-bid second-price auction in a second respect.

The ascending clock auction makes the local comparison salient. Exiting at q and staying in at q to decide again at $q + 1$ give the bidder different outcomes only if the last rival exits at q . In that case staying wins at price q and exiting loses. The choice between exiting and staying in at q is the comparison between the bids q and $q + 1$ in Proposition 2, restricted to the contingencies in which they differ. The ascending clock auction asks this comparison at every price in turn, and a bidder who answers each question according to her value follows a run of the local improvement algorithm of Figure 1. The sealed-bid second-price auction asks for a single number and leads the bidder to no comparison.

Since the ascending clock auction is OSP and also makes the local comparison salient, Li's comparison changes both channels at once and cannot separate them. Displaying a clock while keeping the extensive-form mechanism of the sealed-bid second-price auction can separate the two channels.

[Breitmoser and Schweighofer-Kodritsch \(2022\)](#) run such a design. They replicate the design of [Li \(2017\)](#) and add intermediate treatments. In the passive clock treatment, subjects submit a sealed bid and then watch an ascending clock determine the outcome. In the dynamic clock treatment, subjects bid in an ascending clock auction without seeing when rivals exit. The passive clock has the same extensive-form mechanism as the sealed-bid second-price auction and is not OSP. Obvious strategy-proofness on its own therefore predicts no difference between the passive clock and the sealed-bid second-price auction, and the same holds for every worst-best comparison concept. The passive clock nevertheless makes the comparison between neighboring bids salient, since it resolves the bid price by price.

The passive clock removes about half of the difference in deviations from truthful bidding between the sealed-bid second-price auction and the ascending clock auction in which bidders see their rivals exit. Dynamic bidding, the feature that makes the ascending clock auction OSP, adds a further twelve percent, which is not statistically significant.

One finding in [Breitmoser and Schweighofer-Kodritsch \(2022\)](#) lies outside both channels. Observing rivals' exits during a dynamic clock accounts for almost forty percent of the total effect, although this information is payoff-irrelevant under private values. The effect is significant only when the exits are seen while bidding. I attribute it to learning from the environment, a channel that neither obvious strategy-proofness nor the local improvement model of this paper contains.

This paper's explanation can account for the remaining findings, although Section 7.4 notes that it is not the only explanation that does. Under this explanation, the gain from the sealed-bid second-price auction to the passive clock comes from the local comparison alone. Most of the ascending clock auction's advantage that does not come from observing rivals' exits is already present in the passive clock. The explanation also predicts that the passive clock's effect lags that of the dynamic clock by about one round, which [Breitmoser and Schweighofer-Kodritsch \(2022\)](#) report, as Appendix B.3 explains.

Appendix B discusses further evidence consistent with this paper's explanation. The menu description of [Gonczarowski, Heffetz and Thomas \(2023\)](#) changes what a bidder in the sealed-bid second-price auction is told without changing what she is asked, and it leaves straightforward behavior unchanged.

The ascending clock auction also shows that the verdict of the criterion depends on the neighbor relation. It is local improvement solvable under one-step neighbors on exit prices, the neighbors its dynamic presentation makes salient, and not under one-situation neighbors, as Appendix C.3 shows.

7.2 The static Vickrey auction and the Ausubel auction

The Ausubel auction of [Ausubel \(2004\)](#) is a dynamic implementation of the multi-unit Vickrey rule of Section 4.2. There are Q identical units and n bidders, and a price rises from 0 one grid step at a time. At each price q every bidder states a demand, a number of units no larger than the demand she stated at the previous price. A bidder observes only the current price and whether the auction is still open, the case that [Ausubel \(2004\)](#) calls no bid information. She learns the units she has clinched when the auction ends. When the total demand of the other bidders at price q is $D_{-i}(q)$, the number of units guaranteed to bidder i is $\max\{0, Q - D_{-i}(q)\}$, capped at her own demand. Each unit that becomes guaranteed at price q is clinched and paid for at q , and she keeps it even if her demand later falls below the number of units she has clinched. The auction ends at the first price at which total demand is at most Q . If total demand still exceeds Q at price $K - 1$, the auction ends at K without further clinching, and the units not clinched remain unsold.

A bidder who reduces her demand from x to $x - 1$ at price q withdraws her x -th marginal bid at q . A strategy that reduces demand from ℓ to $\ell - 1$ at price a^ℓ for each ℓ , and never otherwise, has the payoff of the marginal bid vector a in the static Vickrey auction. The Ausubel auction therefore implements the outcomes of the static Vickrey auction, with ϕ sending each vector to this strategy. As the definition of implementing the outcomes requires, sincere bidding is weakly dominant, since with no bid information every strategy has the payoff of some vector of marginal bids in the static Vickrey auction.²²

The Ausubel auction makes the local comparison salient much as the ascending clock auction does, although at each price it asks for a number of units rather than a yes or a no. Reducing demand from x to $x - 1$ units at price q , instead of holding it, changes the bidder's outcome only if the x -th unit would clinch at q . The choice between reducing and holding demand at q is the comparison of Proposition 3 for unit x at bid q . The Ausubel auction asks this comparison for each unit at each price, and a bidder who answers each according to the marginal value of the unit follows a run of the local improvement algorithm. The static Vickrey auction asks for a vector of marginal bids in one move, so the prediction is that the Ausubel auction produces more dominant-strategy play.²³ [Kagel and Levin \(2009\)](#) compare the static Vickrey auction with two Ausubel auctions, one with and one without public dropout information. Bidding is significantly closer to sincere bidding in the version with dropout information than in either of the other two. In that version, however, sincere bidding is an ex post equilibrium rather than a dominant strategy. The gain resembles the effect of observed exits that Section 7.1 attributes to learning from the environment. A passive Ausubel auction, in which marginal bid vectors are submitted and then resolved price by price, would separate the salience of the local comparison from dynamic bidding. Appendix C.4 proposes a further design, a multi-account clock that asks the comparisons of Proposition 3 one unit at a time. Neither it nor the Ausubel auction is OSP for a bidder who demands two or more units, so comparing the two would test the salience of the local comparison with obvious strategy-proofness held fixed.

7.3 Serial dictatorship and dynamic matching mechanisms

The dynamic serial dictatorship is the following extensive-form mechanism. The agents are called one at a time in the priority order, and each agent, when called, observes the set of objects with remaining capacity and picks one of them, which she receives. It implements the outcomes of serial dictatorship with list submission of Section 4.3. In the random serial

²²When bidders observe the others' bids, sincere bidding is an ex post equilibrium, and that version lies outside the class to which the prediction applies.

²³[Nagel and Saitto \(forthcoming\)](#) show that the Ausubel auction is the simplest implementation of the Vickrey outcome under their measure.

dictatorship treatments of [Li \(2017\)](#), the dynamic serial dictatorship produces the dominant-strategy outcome in 93 percent of games, against 64 percent for list submission. [Li \(2017\)](#) explains the difference by obvious strategy-proofness. The dynamic serial dictatorship is OSP, and in fact SOSp in the sense of [Pycia and Troyan \(2023\)](#), since the object the agent takes is the object she receives, while list submission is not OSP.

The dynamic serial dictatorship also makes the local comparison salient, and list submission does not. Suppose the agent is called and compares taking a now with taking b now, with the rest of her strategy fixed. In doing so she compares the list that has a just before b against the list that has b just before a . These are two one-exchange neighbors, and by [Proposition 4](#) they lead to different outcomes only where one gives her a and the other gives her b . The most common error under list submission, the exchange of the two highest entries, is a single adjacent exchange away from the truthful list. This exchange seems to be the error of a participant who has not compared the two lists. As with the ascending clock auction, Li's comparison changes obvious strategy-proofness and the salience of the local comparison at once. The two explanations are separated only by mechanisms that change one without the other.

[Bó and Hakimov \(2024\)](#) provide such mechanisms. In a pick-an-object mechanism, every agent is first shown a menu of objects and picks one. A pick is tentative. An agent whose pick is rejected is shown a new menu, a strict subset of her previous one without the rejected pick, and picks again. Each agent observes only her own menus and picks, and when the procedure ends she receives the last object she picked. The recommended strategy picks the most preferred object on each menu. Since a pick can be rejected, the pick-an-object mechanisms discussed here are not OSP.²⁴ Pick-an-object implementations exist for rules with no OSP implementation at all, including top trading cycles under cyclic priorities and deferred acceptance under general priorities.

For the pick-an-object serial dictatorship the recommended strategy is dominant, and the mechanism implements the outcomes of serial dictatorship with list submission, as [Appendix C.5](#) shows. In the pick-an-object top trading cycles the menus of other agents can depend on an agent's picks, and I treat that mechanism, like iterative deferred acceptance in [Appendix B.4](#), as partial evidence.

The local comparison is more salient in a pick-an-object mechanism than under list submission. Each menu asks the agent for one object, and choosing a over b on a menu, with the rest of her strategy fixed, compares the one-exchange neighbors of [Proposition 4](#). The prediction is therefore that the pick-an-object serial dictatorship produces more dominant-strategy play than

²⁴[Bó and Hakimov \(2024\)](#) show that an individually rational rule has a pick-an-object implementation in perfect ex post equilibrium if and only if the rule is strategy-proof and satisfies a condition they call monotonic discoverability.

serial dictatorship with list submission. Obvious strategy-proofness on its own makes no such prediction, since the pick-an-object serial dictatorship is not OSP.

Bó and Hakimov (2024) find gains that vary in size and in the aspect of behavior they show up in. Under top trading cycles the rate of truthful play rises from 35 percent under list submission to 55 percent under pick-an-object with cyclic priorities, and from 56 to 67 percent with acyclic priorities. Both differences are significant. Under serial dictatorship the rate rises from 72 to 75 percent, which is not significant. But the objects that agents receive are significantly better under pick-an-object than under list submission, measured by the average rank of the assigned object in the agent’s true preference. This suggests that the errors that remain under pick-an-object are more often ones that do not change the outcome. Under the obviously strategy-proof dynamic mechanisms, which have both channels, truthful play is 98 percent for serial dictatorship and 77 percent for top trading cycles with acyclic priorities. Yet the assigned objects are no better than under pick-an-object.

A dynamic mechanism asks for one choice at a time. In each dynamic mechanism of this section, the choice at a move contains a comparison between two neighbors under the neighbor relation of Section 4. These mechanisms therefore make the local comparison salient, one comparison at a time. Making the comparison salient should help when the mechanism is local improvement solvable under those neighbors, so that each comparison has a sign that leads toward the dominant strategy. Dynamic play does not appear to be required for the effect of salience, since the passive clock makes the comparison salient without it. Obvious strategy-proofness is a separate channel. It does not appear to be required for the effect, as the passive clock and, in part, the pick-an-object serial dictatorship and iterative deferred acceptance suggest. Nor does obvious strategy-proofness seem sufficient for dynamic implementation to help, as dynamic Sudoku in Section 6.3 suggests.

7.4 What the explanation does not account for

Table B.1 in Appendix B collects the comparisons discussed in this section, and several of them fit this paper’s explanation only in part or not at all. The passive clock of Breitmoser and Schweighofer-Kodritsch (2022) shows subjects the clock after they have bid. Its effect could therefore come from learning from the feedback of watching their own bid resolved, and not only from the salience of the comparison between neighboring bids.²⁵ The lag of about one

²⁵Breitmoser and Schweighofer-Kodritsch (2022) describe the clock display as inducing a reasoning process that runs through the bids in ascending order, which is close to this paper’s explanation. They also stress that subjects discover dominance by learning from feedback.

round fits both explanations, so the passive clock separates the salience of the local comparison from obvious strategy-proofness without separating salience from learning.

Other findings go against this paper’s explanation or lie outside it. Observing rivals’ exits in the ascending clock auction, and observing dropouts in the Ausubel auction of [Kagel and Levin \(2009\)](#), move bidding closer to the recommended strategy. Yet neither changes obvious strategy-proofness or the salience of the local comparison. Under the pick-an-object serial dictatorship, truthful play does not rise significantly, and the gain appears only in the objects that agents receive.

The rationing procedure of [Brown, Stephenson and Velez \(2026\)](#) makes the local comparison salient and raises the rate of actions consistent with the dominant strategy. The obviously strategy-proof procedure for the uniform rule that they test asks each participant, step by step, whether to stop or to let her allocation move one more unit. It nevertheless produces the outcome of the uniform rule no more often than the direct mechanism, which the authors attribute to mistakes that accumulate over the several moves the procedure requires.²⁶

Finally, the criterion is silent on the direction of errors. Under one-step neighbors the only locally unimprovable bids in the sealed-bid second-price auction are the truthful ones, so the criterion does not predict that subjects overbid, as [Kagel, Harstad and Levin \(1987\)](#) find, rather than underbid. This paper’s explanation is a claim about which of two mechanisms produces more dominant-strategy play. It seems to fit most of the comparisons in Table [B.1](#), but it is not a model of the errors that remain.

8 Design recommendations and conclusion

Strategy-proof mechanisms can be simple in a second way that the worst-best comparison concepts do not measure, since these concepts do not fully use the structure of the strategy set. In some mechanisms, every chain of moves to an improving neighbor ends at a dominant strategy. Local improvement solvability formalizes this property, and strategy-proofness is its weakest version. Many classical strategy-proof mechanisms are local improvement solvable under their natural neighbor relations, in a rich environment for the matching mechanisms, and each of their comparisons reduces to one case. In extensive-form mechanisms and games without duplicate dominant strategies, the property under one-situation neighbors is strict action-proofness, and it neither implies nor is implied by obvious strategy-proofness. Under each game’s natural neighbor relation, the criterion accounts for part of what the worst-best comparison concepts miss in chess, Sudoku, and the pick-a-number game. The evidence, above

²⁶A mechanism with pre-play feedback, which is not OSP, performs better.

all the passive clock, suggests that the local comparison helps participants when a mechanism makes it salient, even with obvious strategy-proofness held fixed.

Two design recommendations follow. The first is to choose, among mechanisms that implement the same outcomes, one that is local improvement solvable under its natural neighbor relation, the relation that the presentation of its strategy set defines. Without this property the comparison between neighbors can stall before reaching the dominant strategy, as it does in chess and Sudoku. Strong obvious strategy-proofness does not guarantee local improvement solvability under the natural neighbor relation, since static Sudoku is SOSP and is not local improvement solvable under one-cell neighbors.

The second recommendation is to make the local comparison salient, so that participants see the path when a mechanism has one. Two presentations do this, and the evidence for the first seems stronger.

First, where submissions must be simultaneous or sealed, resolve them step by step while the participant observes, as the passive clock of [Breitmoser and Schweighofer-Kodritsch \(2022\)](#) does for the sealed-bid second-price auction. A passive presentation leaves the extensive-form mechanism, and therefore its incentive properties, unchanged.²⁷ The evidence for the passive presentation separates the salience of the local comparison from obvious strategy-proofness, although not from learning from feedback.

Second, implement the rule dynamically where the setting permits. The ascending clock auction, the dynamic serial dictatorship, pick-an-object mechanisms, and iterative deferred acceptance each ask the participant for one choice at a time, and each choice contains a comparison between neighboring strategies. [Li \(2017\)](#), [Breitmoser and Schweighofer-Kodritsch \(2022\)](#), [Bó and Hakimov \(2020\)](#), and [Bó and Hakimov \(2024\)](#) find corresponding gains, including for rules that have no obviously strategy-proof implementation.²⁸ The rationing procedure of [Brown, Stephenson and Velez \(2026\)](#), which also asks one comparison at a time and shows no gain in outcomes, is a reminder that a dynamic implementation does not guarantee better outcomes.

The criterion captures only one dimension of simplicity. It takes the neighbor relation as given and does not say which relation a participant actually uses. It records whether every run of local comparisons ends at the dominant strategy, not how hard each comparison is to evaluate, which is what the worst-best comparison concepts measure. And it can hold only

²⁷A sealed-bid online auction, for example, can display the price rising after the bids are in, and a sealed-bid procurement auction can display it falling. A related static presentation can elicit willingness to pay by the Becker-DeGroot-Marschak procedure with a list that asks whether to buy at each price, rather than asking for a single number.

²⁸Appendix C.4 proposes a design that would test the second recommendation in multi-unit auctions with obvious strategy-proofness held fixed.

in strategy-proof mechanisms. The first and the last of these limits lead to two of the open questions below.

Three questions are open. The first concerns mechanisms that are not strategy-proof, where the criterion cannot hold. In the strategy-proof mechanisms of Section 4, the comparison between two neighbors has a sign that does not depend on a belief about the others. Iterative deferred acceptance already lies outside the class of strategy-proof mechanisms. Its recommended strategy is an equilibrium strategy rather than a dominant one, and yet asking for one choice at a time appears to bring a similar gain there. Do comparisons between neighbors lead a participant to a best response given her belief about the others? Which condition on the neighbor relation and the outcome function makes every run of the local improvement algorithm end at such a best response? And how does that condition relate to the equilibrium concepts under which mechanisms such as iterative deferred acceptance are analyzed?

The second question concerns design, and it stays within the strategy-proof rules. Local improvement solvability under a natural neighbor relation is stronger than strategy-proofness, so a rule implemented by a mechanism that satisfies it must be strategy-proof. Which strategy-proof rules have such an implementation? [Yang and Dworczak \(2026\)](#) impose a requirement of this kind on menus. They call an ordered menu searchable if, for every type, every local maximum of utility over the menu is global, up to a condition on ties. They show that searchability restricts multidimensional screening. For a finite menu, the condition on local maxima is local improvement solvability for every type under one-step neighbors on the menu's order. This paper may therefore serve as a micro-foundation for their approach. It argues that comparing neighbors can simplify the search for the best strategy, and Section 7 suggests when participants make such comparisons. The question remains open for rules with several participants.

The third question concerns a model of salience. A model of what a description or a presentation must do to induce a given neighbor relation, in the spirit of the descriptions of [Gonczarowski, Heffetz and Thomas \(2023\)](#), would make the prediction sharper.²⁹ What a mechanism asks a participant to compare, and not only what the mechanism is, appears to shape whether she finds the dominant strategy.

²⁹The finding of [Breitmoser and Schweighofer-Kodritsch \(2022\)](#) that observing rivals' exits accounts for a large share of the effect of the ascending clock auction lies outside such a model. Section 7.1 attributes the effect of observed exits to learning from the environment. Combining the comparison of neighbors with a model of learning is a further direction.

A Proofs

A.1 Proof of Proposition 1

Part (a). A locally unimprovable strategy s_i under N'_i is a strategy that no strategy in $N'_i(s_i)$ improves. Since $N_i(s_i) \subseteq N'_i(s_i)$, no strategy in $N_i(s_i)$ improves it either, so it is locally unimprovable under N_i . Every locally unimprovable strategy under N'_i is therefore a locally unimprovable strategy under N_i , all of which are payoff-equivalent by local improvement solvability under N_i . Hence Γ is local improvement solvable for θ_i under N'_i .

Part (b). Under the complete neighbor relation, a strategy is locally unimprovable if no strategy improves it. If Γ is local improvement solvable under the complete neighbor relation, then θ_i has a dominant strategy by Theorem 1(a). Conversely, let s_i^* be dominant and let p be locally unimprovable. Then s_i^* gives at least the same utility as p in every contingency. If it gave a higher utility in some contingency, it would improve p , and p would not be locally unimprovable. So p is payoff-equivalent to s_i^* , all locally unimprovable strategies are payoff-equivalent, and Γ is local improvement solvable for θ_i under the complete neighbor relation. $\square$

A.2 The measure of Nagel and Saitto

I state the measure of Nagel and Saitto (forthcoming) for a single strategic player i in an extensive-form mechanism, the other players and nature being folded into the contingencies. For a situation $h \in H_i$ and a strategy $\rho \in S(h)_i$, let $Z(\rho, h) = \{u_i(\rho, s_{-i}; \theta_i) : s_{-i} \in S(h)_{-i}\}$ be the set of payoffs the player can receive when play reaches h and she follows ρ from h on. For an action $a \in A(h)$, let $Z(a, h)$ be the union of $Z(\rho, h)$ over the strategies $\rho \in S(h)_i$ with $\rho_h = a$, the payoffs that can follow a whatever she does later.

Definition 16 (Evident reduction): *At a situation h , action a is evidently never better than a' if $\max Z(a, h) \leq \min Z(a', h)$. The evident reduction at h is the set of actions that survive iterated removal, one action at a time, of an action evidently never better than some other surviving action.*

Definition 17 (Essential frame and complexity): *An essential frame merges two consecutive situations of the player whenever the information revealed between them excludes no payoff possible under her dominant continuation. An action at the merged situation specifies her actions at both. The complexity of an extensive-form mechanism is the largest number of actions in an evident reduction of its essential frame, over the players, their types, and their situations.*

The complexity is 1 for the simplest mechanisms, and in the most complex static mechanisms it equals the number of strategies. Nagel and Saitto (forthcoming) define an essential frame directly, as a strategy-proof mechanism in which new relevant information is revealed to the agent between every two of her consecutive situations and whose tree is minimal. They define the complexity of any strategy-proof mechanism through an evident reduction of the player's continuation behaviors, which in an essential frame becomes the comparison of actions in Definition 16. They also show that every strategy-proof mechanism has an essential frame with the same outcomes and weakly lower complexity. The two definitions above simplify their construction.

A.3 Proof of Proposition 3

The residual prices are functions of the rivals' bids, weakly increasing in ℓ , and do not depend on the agent's strategy. By the allocation rule, unit ℓ is won if and only if $a^\ell > p_\ell$. The strategy is weakly decreasing and the residual prices are weakly increasing, so the units the agent wins are units 1 through x for some x , and the payoff is (2). Changing a^ℓ to $a^\ell + 1$ changes the term for unit ℓ only, and only in contingencies with $p_\ell = a^\ell$, where it changes the term from 0 to $\theta^\ell - a^\ell$. Some profile of rival bids gives $p_\ell = a^\ell$, for example one in which the Q highest rival bids all equal a^ℓ . Three consequences follow, provided the changed strategy is weakly decreasing. The change from a^ℓ to $a^\ell + 1$ is a local improvement when $a^\ell < \theta^\ell$. The change from a^ℓ to $a^\ell - 1$ is a local improvement when $a^\ell - 1 > \theta^\ell$. A change between two bids in $T(\theta^\ell)$, which occurs only when θ^ℓ is a grid point and the bids are θ^ℓ and $\theta^\ell + 1$, changes no payoff.

Write $\underline{t}^\ell$ and $\bar{t}^\ell$ for the smallest and largest elements of $T(\theta^\ell)$. Since $\theta^\ell > \theta^{\ell+1}$, every element of $T(\theta^\ell)$ is at least every element of $T(\theta^{\ell+1})$. If $\theta^{\ell+1}$ is a grid point, then $\bar{t}^{\ell+1} = \theta^{\ell+1} + 1 \leq \underline{t}^\ell$, because $\underline{t}^\ell$ is a grid point at or above $\theta^\ell > \theta^{\ell+1}$. Otherwise $\bar{t}^{\ell+1}$ is the least grid point above $\theta^{\ell+1}$, which is at most $\underline{t}^\ell$.

Feasibility. Let a be a weakly decreasing strategy that is not truthful. Suppose first that some unit has $a^\ell > \bar{t}^\ell$, and take the largest such ℓ . Then $a^\ell - 1 \geq \bar{t}^\ell \geq \theta^\ell$, and $a^\ell - 1 > \theta^\ell$, because $a^\ell - 1 = \theta^\ell$ would put a^ℓ in $T(\theta^\ell)$. If $\ell < L$, then $a^{\ell+1} \leq \bar{t}^{\ell+1} \leq \underline{t}^\ell \leq a^\ell - 1$. So lowering a^ℓ by one keeps the strategy weakly decreasing, and the change is a local improvement. Suppose instead that every unit has $a^\ell \leq \bar{t}^\ell$. Then some unit has $a^\ell < \underline{t}^\ell$, so $a^\ell < \theta^\ell$. Take the smallest such ℓ . If $\ell > 1$, then $a^{\ell-1} \geq \underline{t}^{\ell-1} \geq \bar{t}^\ell \geq \underline{t}^\ell \geq a^\ell + 1$. So raising a^ℓ by one keeps the strategy weakly decreasing, and the change is a local improvement. Hence no strategy that is not truthful is locally unimprovable.

Every truthful strategy is locally unimprovable. Each of its neighbors either moves some a^ℓ to the other element of $T(\theta^\ell)$, which changes no payoff, or moves a^ℓ outside $T(\theta^\ell)$. By the first paragraph, the second kind of move lowers the payoff in the contingencies where p_ℓ equals the lower of the two bids and changes it nowhere else. All truthful strategies give the same utility in every contingency. Each unit's term depends on a^ℓ only through whether $a^\ell > p_\ell$, and the two elements of $T(\theta^\ell)$, when there are two, differ in this only at $p_\ell = \theta^\ell$, where the term is 0 either way. $\square$

A.4 Proof of Proposition 4

Part (a). Fix a type θ and a strategy $s_i \neq \theta$. Some adjacent pair of s_i is ordered differently from θ , since otherwise every adjacent pair of s_i would be ordered as in θ and s_i would equal θ . Let s'_i be the strategy obtained from s_i by exchanging such a pair o, o' , where θ ranks o above o' and s'_i lists o first. By the adjacent-swap property applied to s'_i and s_i , for every s_{-i} either $f(s'_i, s_{-i}) = f(s_i, s_{-i})$ or $f(s'_i, s_{-i}) = o$ and $f(s_i, s_{-i}) = o'$. Hence s'_i gives at least the same utility as s_i in every contingency and a higher utility in every contingency where the pair is decisive. Richness provides such a contingency. So s'_i locally improves s_i , and no strategy other than θ is locally unimprovable. The strategy θ is locally unimprovable. Each of its neighbors exchanges a pair into the order opposite to θ and, by the same argument with the roles exchanged, gives at most the same utility in every contingency. Every run therefore ends at θ , and θ is dominant by Theorem 1.

Part (b). Fix a contingency s_{-i} and two strategies s_i and s'_i that differ by an adjacent exchange of o and o' , with s_i listing o first. Let $x = f(s_i, s_{-i})$ and $y = f(s'_i, s_{-i})$, and suppose $x \neq y$. The type whose true ranking is s_i exists by the full domain. Strategy-proofness for that type says that she does at least as well with s_i as with s'_i , so s_i ranks x above y . Strategy-proofness for the type whose true ranking is s'_i says that s'_i ranks y above x . The two lists order every pair of objects in the same way except the pair $\{o, o'\}$, so $\{x, y\} = \{o, o'\}$, and since s_i lists o first, $x = o$ and $y = o'$. The only difference an adjacent exchange can make to the outcome is to exchange the two objects that were exchanged in the list. $\square$

A.5 Richness in the three matching mechanisms

This subsection verifies the example of a rich environment given after Corollary 1. Fix a list and an adjacent pair at positions j and $j + 1$. In serial dictatorship, the pair is decisive in a contingency in which the agent moves at position j or later and her predecessors exhaust the capacities of exactly the $j - 1$ objects above the pair. In top trading cycles, the pair is decisive

in a contingency in which two conditions hold. First, at each of the $j - 1$ objects above the pair, the agents with the highest priority, as many as its capacity, are other agents who rank that object first. Second, the agent has top priority at both objects of the pair. Until its capacity is exhausted, each object above the pair points to one of these agents, who points back to it. The agent points to such an object while one has capacity, so she is in no cycle until all of them are exhausted. Both objects of the pair point to her while she remains, so neither loses capacity before she is assigned, and she then receives the one her list ranks first. In deferred acceptance, the pair is decisive in a contingency in which two conditions hold. First, the $j - 1$ schools above the pair fill their capacities with students who rank them first and have higher priority there than the agent. Second, the agent has top priority at both schools of the pair. From the first round on, each school above the pair is then full with students of higher priority than the agent, so it rejects her. Neither school of the pair rejects her, so she is assigned the one her list ranks first. Suppose the agent's priority is part of the contingency, every profile of priorities and others' strategies is possible, and there are enough other agents to exhaust the capacities of the objects above the pair. Then each of these contingencies exists. $\square$

When the agent knows that some adjacent pair on her list cannot affect her outcome, the two strategies that exchange the pair are payoff-equivalent, and the exchange is not a local improvement. Suppose, for example, that the agent knows that she is second in a serial dictatorship over four objects, one unit each, and her true ranking is $abcd$. The list $acdb$ has no local improvement. Exchanging d and b changes no outcome, and the other two adjacent exchanges lower her payoff. Yet the list gives her c instead of b when the first dictator takes a . Richness rules out such a list.

A.6 Dominance at a situation

This subsection proves the equivalence of (1) and (D) in Definition 11. Condition (1) implies (D), since (D) restricts (1) to some strategies and contingencies. For the converse, suppose a strategy reaches a different terminal node from s_i^* in some contingency. Then the two strategies' action sequences in that contingency first differ at some situation $h \in H_i$. Both strategies are therefore in $S(h)_i$, the contingency reaches h , and the strategies specify different actions at h . Hence (D) at h gives the inequality in (1) for that contingency. In a contingency in which the two strategies reach the same terminal node, they give the same utility. $\square$

A.7 Proof of Theorem 2

Two strategies that specify the same action at every situation both of them reach lead to the same play path in every contingency, and are therefore payoff-equivalent.

From strict action dominance to local improvement solvability. Let s_i^* be strictly action dominant. It is dominant by Proposition 5(a), hence locally unimprovable. Let p be locally unimprovable and suppose p and s_i^* both reach some h at which they differ. Then $p \in S(h)_i$, and by strict action dominance $p[h \leftarrow s_h^*]$ locally improves p , a contradiction. So every locally unimprovable strategy agrees with s_i^* wherever both reach and is payoff-equivalent to it. A strictly action dominant strategy also rules out duplicates. Suppose a dominant strategy p differed from s_i^* at a situation h that both reach. Strict action dominance at h with $r = p$ would then make $p[h \leftarrow s_h^*]$ strictly better than p in some contingency, which contradicts the dominance of p . So every dominant strategy agrees with s_i^* wherever both reach. If two dominant strategies differed at a situation g that both reach, s_i^* could not reach g , since it would agree with both there. It would then differ from both at the situation where it leaves the path to g , which all three reach, a contradiction.

From local improvement solvability to strict action dominance. Let Γ be local improvement solvable under one-situation neighbors with no duplicate dominant strategies for θ_i , and let s_i^* be a dominant strategy, which exists by Lemma 1 and Theorem 1(a). Fix a situation h that s_i^* reaches, a strategy $r \in S(h)_i$, and an action $a' \neq s_h^*$, and suppose strict action dominance fails for this situation, strategy, and action. Then either $u_i(r[h \leftarrow a'], s_{-i}^0; \theta_i) > u_i(r[h \leftarrow s_h^*], s_{-i}^0; \theta_i)$ for some s_{-i}^0 that reaches h , or $r[h \leftarrow a']$ and $r[h \leftarrow s_h^*]$ are payoff-equivalent. Let $q = r[h \leftarrow a']$, and call the situations h' of player i with $S(h')_h = \{s_h^*\}$ the recommended subtree. Run the local improvement algorithm from q with changes restricted to situations outside the recommended subtree, and let p be the strategy at which it stops. Along the run the utility weakly increases in every contingency, so p does at least as well as q everywhere.

The run never changes the action at h to s_h^* . Such a change, from a strategy p' to $p'[h \leftarrow s_h^*]$, alters no outcome if p' does not reach h . If $p' \in S(h)_i$, it alters play only in contingencies that reach h . In those contingencies, $p'[h \leftarrow s_h^*]$ plays as $r[h \leftarrow s_h^*]$, because the actions in the recommended subtree are still those of r . In the first case the changed strategy is worse than q , and hence than p' , at s_{-i}^0 . In the second case it gives the utility of q in every contingency that reaches h and that of p' in every other, hence at most the utility of p' everywhere. So the change is not a local improvement. Hence $p_h \neq s_h^*$, no situation in the recommended subtree is reached under p , and the changes the run did not consider alter no outcome. So p is locally unimprovable, hence dominant by Theorem 1. Since there are no duplicates, p agrees with s_i^* at every situation both reach. If p reaches h , this agreement fails at h . If not, then at the earliest

situation g preceding h at which p leaves the path to h , both strategies reach g and differ there. Either way there is a contradiction, so s_i^* is strictly action dominant. $\square$

A.8 Proof of Proposition 5

Part (a). Fix s_i^* , a situation h it reaches, and $a' \neq s_h^*$, with minima and maxima ranging as in Definitions 12 and 14. $(SOD) \Rightarrow (OD)$ holds because s_i^* is a continuation of s_h^* and every deviation s'_i at h with $s'_h = a'$ is a continuation of a' . Both this implication and $(OD) \Rightarrow (D)$ use the fact that a minimum is at most and a maximum at least each element of its set. $(SOD) \Rightarrow (AD)$ holds because $r[h \leftarrow s_h^*]$ and $r[h \leftarrow a']$ are continuations of s_h^* and a' . Hence, in every contingency that reaches h , the first pays at least the minimum in (SOD) and the second at most the maximum, while in other contingencies they coincide.

$(AD) \Rightarrow (D)$. Let s'_i deviate from s_i^* at h , so that $s'_h = a'$. Define r by $r_{h''} = s'_{h''}$ at every situation h'' that follows h through a' and $r_{h''} = s_{h''}^*$ at every other situation. Then $r \in S(h)_i$, since r agrees with s_i^* at the situations preceding h . Fix $s_{-i} \in S(h)_{-i}$. The play of $(r[h \leftarrow s_h^*], s_{-i})$ passes through h , and after h it meets only situations of player i that follow h through s_h^* . At all of these, and at the situations preceding h , r agrees with s_i^* , so the play reaches the terminal node of (s_i^*, s_{-i}) . In the same way $(r[h \leftarrow a'], s_{-i})$ reaches the terminal node of (s'_i, s_{-i}) , since s'_i agrees with s_i^* before h . Condition (AD) at h therefore gives $u_i(s_i^*, s_{-i}; \theta_i) \geq u_i(s'_i, s_{-i}; \theta_i)$, which is (D) at h . Since h and s'_i were arbitrary, (D) holds at every situation that s_i^* reaches, and by Appendix A.6, s_i^* is dominant.

Parts (b) and (c). Section 5.3 shows that the game of Figure 3 is OSP and not AP and that the sealed-bid second-price auction is AP and not OSP. The other two examples are the games of Figure A.1. In each tree the other player j moves first and chooses the contingency, c_1 or c_2 . The agent i does not observe this choice, so each of her situations contains two nodes, joined by a dashed line, and each terminal node shows her payoff only. Strategies are written as in Section 5.2, and a pair such as $(10, 3)$ lists the agent's payoffs at c_1 and c_2 .

The game on the left of Figure A.1 is SP and neither OSP nor AP. The strategy (L, x) pays $(10, 3)$ and is dominant. It is not obviously dominant, since its worst payoff, 3, is below the best payoff of R , 5. It is not action dominant at the first situation, since (L, y) pays $(0, 0)$ and (R, y) pays $(5, 2)$.³⁰

The game on the right of Figure A.1 proves (c). The recommended strategy (a^*, b^*) pays $(10, 10)$, the best payoff in each contingency, so it is obviously dominant. Changing either of its actions weakly lowers the payoff in each contingency, whatever the other action is, so the

³⁰Under one-situation neighbors the game is not local improvement solvable either. The strategy (R, y) has no local improvement, since changing R to L gives $(0, 0)$ and changing y changes nothing at an unreached situation.




Figure A.1: The games of Proposition 5(b), left, and Proposition 5(c), right.

recommended strategy is action dominant. It is not strongly obviously dominant, since the worst payoff after a^* , over the agent's continuations and the contingencies, is 1, below the best payoff after a' , 9. Strong obvious dominance compares the minimum over own continuations and contingencies jointly against the maximum over both jointly. Action dominance fixes the contingency, obvious dominance fixes the continuation, and neither compares a bad continuation in one contingency against the deviation in another. $\square$

A.9 Proof of Proposition 6(c)

Step one. Let ϕ and σ be as in Definition 15, so that $\sigma_i = \phi_i(K)$ for each player. Let h be the first situation of player i on the play of $(\phi_i(a), \phi_j(b))$ at which $\phi_i(a)$ and $\phi_i(a')$ choose different actions, where $a < a'$ without loss of generality. It exists, since the profiles $(\phi_i(a), \phi_j(b))$ and $(\phi_i(a'), \phi_j(b))$ give the outcomes (a, b) and (a', b) and differ only in the strategy of i , so their plays first differ at a situation of i . The two strategies agree before h , so the play of $(\phi_i(a'), \phi_j(b))$ also reaches h . Suppose that some $\phi_j(b'')$ does not reach h . For $c \in \{a, a'\}$, let τ_c agree with $\phi_i(c)$ at h and at every situation that follows h , and with the other of the two strategies elsewhere. Against $\phi_j(b)$ the strategy τ_c plays as $\phi_i(c)$. Against $\phi_j(b'')$ the plays of both strategies avoid h and every situation that follows it, so τ_c plays as the other strategy. If $b'' = K$, then against $\tau_{a'}$ the strategy $\phi_j(b)$ gives j the payoff $b + \lambda a' \geq \lambda a'$, and $\phi_j(K)$ gives her $K + \lambda a$. If $b = K$, then against τ_a the strategy $\phi_j(b'')$ gives j at least $\lambda a'$, and $\phi_j(K)$ gives her $K + \lambda a$. Since $\lambda a' - K - \lambda a \geq \lambda - K > 0$, in either case $\phi_j(K)$ is not dominant, which contradicts strategy-proofness under σ . If neither b nor b'' is K , the first case applies to K in place of b'' if $\phi_j(K)$ does not reach h , and the second to K in place of b otherwise. Hence every $\phi_j(b')$ reaches h .

Step two. Fix $m < K$. Against $\phi_j(0)$ the strategies $\phi_i(K)$ and $\phi_i(m)$ give the outcomes $(K, 0)$ and $(m, 0)$. Let h_0 be the first situation of player i on the play of $(\phi_i(K), \phi_j(0))$ at which

the two strategies differ. By step one, $\phi_j(K)$ reaches h_0 . Over the contingencies that reach h_0 , the minimum payoff of $\phi_i(K)$ is at most K , at $\phi_j(0)$, and the maximum payoff of $\phi_i(m)$ is at least $m + \lambda K > K$, at $\phi_j(K)$. So (OD) fails at h_0 . Every dominant strategy of i pays what $\phi_i(K)$ pays in every contingency, $K + \lambda b$ against $\phi_j(b)$, and so gives the outcome (K, b) , since $\lambda > K$. Replacing $\phi_i(K)$ by it preserves Definition 15, so the same argument applies to it. Hence no extensive-form game that implements the outcomes of the static game is OSP.

Step three. Now suppose that, for every a , the strategy $\phi_i(a)$ gives i the number a against every strategy of j . Let h be a situation of i that $\phi_i(x)$ and $\phi_i(y)$ reach, with $x < y$, and suppose that some strategy of j does not reach h . Then there are two strategies of j , one of them $\phi_j(K)$, of which exactly one reaches h . On the play of $\phi_i(x)$ against the one that reaches h , let g be the first situation of i that the play of $\phi_i(x)$ against the other does not reach. The strategies $\phi_i(x)$ and $\phi_i(y)$ agree before g , and against every strategy of j they give different numbers. So the play of $\phi_i(x)$ against the other strategy reaches a situation $g' \neq g$ of i after the same actions of i , and both strategies reach g and g' . No play reaches both g and g' , as neither precedes the other, so neither strategy of j reaches both, and each reaches one of them. Let τ agree with $\phi_i(y)$ at the one of g and g' that $\phi_j(K)$ does not reach and at every situation that follows it, and with $\phi_i(x)$ elsewhere. Against the strategy of j that reaches it, τ gives i the number y , so that strategy of j receives at least λy . Against $\phi_j(K)$, the strategy τ gives i the number x , so $\phi_j(K)$ receives at most $K + \lambda x$. Since $\lambda y - K - \lambda x \geq \lambda - K > 0$, this contradicts strategy-proofness. Hence every strategy of j reaches every situation of i that two of the strategies $\phi_i(a)$ reach, and nothing about j is revealed between two consecutive such situations. In the construction of Nagel and Saitto (forthcoming), two consecutive situations of player i are merged unless the information revealed between them excludes some payoff possible under her dominant continuation. Nothing is excluded here, whether or not $\phi_i(K)$ reaches the two situations. Her first situation is such a situation, as all her strategies reach it, so every play reaches it. Every situation of i before such a situation is again one. So the essential frame merges all these situations into a single situation. Against $\phi_j(0)$, every strategy of i pays at most K , the payoff of the dominant strategy $\phi_i(K)$. So every action at the merged situation has a minimum payoff of at most K . The strategies $\phi_i(a)$ take $K + 1$ different actions there, since any two of them first differ at a situation that both reach. The action of $\phi_i(a)$ can be followed by $a + \lambda K > K$, against $\phi_j(K)$. Hence none of these actions is evidently never better than another action, all of them are in the evident reduction, and the complexity is at least $K + 1$. $\square$

A.10 The chess argument

The strategy of Section 6.2 makes its other moves so that no other change at one position improves it. The winning strategy is chosen so that, whenever a position it reaches offers a move after which every continuation of White wins against all play of Black, it makes such a move. No position on the path to p , p included, then offers such a move. At a position before p the winning strategy would make it, and every continuation would then win at p , m included. At p the move would be w , and the move at q could not be losing. At the other positions that the winning strategy reaches, the strategy plays as the winning strategy. After every other move that the winning strategy does not make, the strategy plays so that Black can prevent a win, where this is possible. After m it plays a continuation that wins against some error of Black and that no change at one position improves. Such a continuation exists by Lemma 1, since a run of the local improvement algorithm after m , started from a continuation that wins against some error of Black, never lowers a payoff. A change after m is then not an improvement. Nor is a change at any other reached position not on the path to p , since the strategy plays as the winning strategy there and wins against all play of Black through it. A change to a new move at p or at a position before p is not an improvement either. Against a strategy of Black that leads to p , errs after m so that the original strategy wins, and prevents a win after the new move, the changed strategy does not win. $\square$

B Additional evidence

This appendix collects the evidence of Section 7 in one table and discusses three groups of studies that the section mentions only briefly or not at all. It also checks richness in Li's random serial dictatorship design.

B.1 Summary of the evidence

Table B.1 collects the experimental comparisons between two mechanisms that produce the same outcomes, or between two descriptions of one mechanism. Each row changes a baseline into a treatment. The second column says whether the change makes the mechanism OSP, and the third says whether it makes the local comparison more salient. Rows marked with a dagger involve a recommended strategy that is an equilibrium strategy rather than a dominant strategy, so they are partial evidence. Section 7 and this appendix discuss the rows in turn.

The rows can be grouped by the channels they change. In the rows where only the salience of the local comparison changes, the gain is large for the passive clock. It is also present for the

pick-an-object top trading cycles and iterative deferred acceptance, although these two rows are partial evidence. For the pick-an-object serial dictatorship, the gain appears in the objects received and not in truthful play. Where only obvious strategy-proofness changes, truthful play rises by 23 percentage points for serial dictatorship and by 10 for top trading cycles, although the objects received are no better. The gain is not significant for the dynamic clock. The ascending clock auction and the dynamic serial dictatorship change both channels and show the largest gains. The rationing procedure of [Brown, Stephenson and Velez \(2026\)](#) also changes both channels, however, and shows no gain in outcomes, although actions consistent with the dominant strategy become more frequent. Section 7.4 takes up this result together with two effects that change neither channel.

B.2 Descriptions of the sealed-bid second-price auction

[Gonczarowski, Heffetz and Thomas \(2023\)](#) change what a bidder in the sealed-bid second-price auction is told without changing what she is asked. Their menu description tells the bidder that the rivals' bids fix a price p and that her own bid determines only whether she buys at that price. This description exposes strategy-proofness without requiring the bidder to reason through the rules. With five bidders, they find no difference in straightforward behavior between the menu description and the traditional description, although the same kind of description raises straightforward behavior in a median-voting mechanism.

This paper's explanation is consistent with the auction result. The menu description asks for a single number, as the traditional description does. It does display a two-element choice, buying at p or not, and this is the decisive case of Proposition 2 stated for a given price. But this choice is a comparison between two outcomes at a price the bidder does not know, not a comparison between two of her own bids. The menu description also does not tell her which bid to move to. Neither description makes the comparison between neighboring bids salient, so the two descriptions leave the bidder with the same task and should perform alike. A null result can have other causes, so the finding supports the explanation less strongly than the passive-clock results do. What seems to separate the passive and dynamic clocks from both descriptions is whether the bidder is led to compare a bid with the next bid up, not what she is told about the rules.

B.3 Further evidence on the ascending clock auction

Further evidence on the ascending clock auction is consistent with this paper's explanation. [Harstad \(2000\)](#) finds that experience with the ascending clock auction partly carries over to

Baseline → treatment	Becomes OSP	Comparison more salient	Finding	Source
<i>Auctions</i>				
Sealed-bid → ascending clock	yes	yes	more dominant-strategy play	Li (2017)
Sealed-bid → passive clock	no	yes	about half of the gap to the clock closed	Breitmoser and Schweighofer-Kodritsch (2022)
Passive clock → dynamic clock	yes	no	a further 12 percent, not significant	Breitmoser and Schweighofer-Kodritsch (2022)
Dynamic clock → dynamic clock showing rivals' exits	no	no	almost 40 percent of the total effect	Breitmoser and Schweighofer-Kodritsch (2022)
Traditional → menu description of the sealed-bid second-price auction	no	no	no difference	Gonczarowski, Heffetz and Thomas (2023)
Ausubel without → with dropout information [†]	no	no	closer to sincere bidding than static Vickrey or Ausubel without dropout information	Kagel and Levin (2009)
<i>Matching</i>				
List submission → dynamic serial dictatorship	yes	yes	64 → 93 percent	Li (2017)
List submission → pick-an-object, serial dictatorship	no	yes	72 → 75 percent, not significant, better objects received	Bó and Hakimov (2024)
Pick-an-object → OSP, serial dictatorship	yes	no	75 → 98 percent, objects received no better	Bó and Hakimov (2024)
List submission → pick-an-object, top trading cycles [†]	no	yes	35 → 55 percent (cyclic), 56 → 67 percent (acyclic)	Bó and Hakimov (2024)
Pick-an-object → OSP, top trading cycles [†]	yes	no	67 → 77 percent (acyclic), objects received no better	Bó and Hakimov (2024)
List submission → iterative deferred acceptance [†]	no	yes	more straightforward play, more stable outcomes	Bó and Hakimov (2020)
<i>Rationing, uniform rule</i>				
Direct → OSP procedure	yes	yes	more actions consistent with the dominant strategy, uniform-rule outcome no more often	Brown, Stephenson and Velez (2026)
Direct → pre-play feedback	no	unclear	more peak strategies and uniform-rule outcomes	Brown, Stephenson and Velez (2026)

Table B.1: Experimental evidence on the two channels, Section 7.

later play in the sealed-bid second-price auction. The carryover is expected if what subjects learn from the ascending clock auction is a way of evaluating bids rather than a property of that auction. [Harstad \(2000\)](#) reads part of it this way, as considering prices one by one, but most of it as an awareness that overbidding leads to losses. [Schotter, Weigelt and Wilson \(1994\)](#) find, in a different game, that presenting a static game as a sequence of moves increases dominant-strategy play. [Georganas, Levin and McGee \(2017\)](#) find that changes in payoffs away from the dominant strategy change behavior sharply in the sealed-bid second-price auction. By contrast, [McGee and Levin \(2019\)](#) find that similar changes alter bidding little in the ascending clock auction. The difference is expected if subjects in the sealed-bid second-price auction evaluate payoffs directly while subjects in the ascending clock auction answer the question the mechanism asks. The Becker-DeGroot-Marschak procedure is a one-bidder second-price auction against a random price. Its price-list version, a table asking whether to sell or keep at each listed price, asks for one comparison at each price instead of a single number. Under the single-number version of the same procedure, [Cason and Plott \(2014\)](#) find that on their first decision five of six subjects misprice a card whose value is known to be two dollars.

[Breitmoser and Schweighofer-Kodritsch \(2022\)](#) also estimate structural models of bidding. In every mechanism that displays a clock, including the static ones, behavior is best explained by a model in which the subject decides at each price whether to continue or stop. Only behavior in the sealed-bid second-price auction without a clock is best explained by a model that evaluates all bids at once.

This paper's explanation also predicts the timing of the passive clock's effect. In the passive clock treatment the clock runs after the bids are submitted, so a subject's first bid comes before she has seen a clock. From the second round on she has watched the clock resolve her own bid price by price. The comparison between exiting at one price and staying to the next is then available when she bids. The passive clock should therefore lag the dynamic clock by about one round and then catch up. Result 3 of [Breitmoser and Schweighofer-Kodritsch \(2022\)](#) reports a lag of approximately one round when dropouts are not shown, and of two rounds when they are. They themselves trace the one-round lag to the clock being shown only after bids are submitted.

B.4 Iterative deferred acceptance

Iterative deferred acceptance provides partial evidence of the kind discussed in Section 7.3. By [Ashlagi and Gonczarowski \(2018\)](#), no stable matching mechanism has an obviously strategy-proof implementation for general priorities of the colleges. In iterative deferred acceptance, students apply to one college at a time and observe the result of each application before the

next. [Bó and Hakimov \(2020\)](#) find that iterative deferred acceptance produces stable outcomes significantly more often than list submission. The reason is that more subjects follow the straightforward strategy, which applies at each round to the most preferred college among those still open to the student. The effect persists when cutoff information is withheld. I read this finding the way I read the ascending clock auction. Each round asks the student for one application. Suppose she compares applying to a now with applying to b now, and keeps the order of her later applications as it was. She then compares two lists that differ by the adjacent exchange of a and b . Iterative deferred acceptance can show her, one round at a time, that the exchange changes her outcome only between a and b .

The evidence is partial for two reasons. First, the straightforward strategy in the iterative mechanisms of [Bó and Hakimov \(2020\)](#) is an equilibrium strategy rather than a dominant strategy, so these mechanisms lie outside the class to which the prediction applies. Second, the iterative designs combine one choice at a time with feedback between rounds. Like the passive clock of [Breitmoser and Schweighofer-Kodritsch \(2022\)](#), they therefore do not separate salience from learning from feedback. The evidence is nonetheless consistent with this paper's explanation.

B.5 Richness in Li's random serial dictatorship design

In the random serial dictatorship treatments of [Li \(2017\)](#), four agents share four objects, one unit each. Before she submits her list, each subject is told her own priority score, an integer from 1 to 10, but not the scores of the others. The others' scores are drawn independently and ties are broken at random, so every position remains possible whatever her score. The environment is therefore rich, and Corollary 1 applies to the design of [Li \(2017\)](#). Her score makes some positions unlikely, however. An agent who moves at position k always receives one of the first k objects on her list, so a pair at positions j and $j + 1$ is decisive for her only when $j \leq k$. The agent who moves last receives the one remaining object whatever her list, and no pair is decisive for her. For a subject with a high score, the pairs low on her list are therefore decisive only in unlikely contingencies. The exchange of the two highest entries, by contrast, can be decisive whenever she does not move last. A list whose first two entries are the two objects she likes best, in the wrong order, therefore has a local improvement, namely the exchange that restores them. Exchanging the two objects she likes best is the most common error that [Li \(2017\)](#) reports under serial dictatorship with list submission.

C Further discussion

This appendix collects discussion that the main text summarizes. It covers the roles of the two requirements in Definitions 5 and 8, the place of other conditions in Table 1, the ascending clock auction under one-situation neighbors, a proposed multi-account clock, and the pick-an-object serial dictatorship.

C.1 The two requirements in the definition

The requirement that a local improvement be strictly better in some contingency and the requirement that all locally unimprovable strategies be payoff-equivalent have different roles. The first requirement makes every run terminate, as Lemma 1 shows. Without it, a run could move back and forth between two payoff-equivalent strategies. The second requirement makes the locally unimprovable strategies dominant. Termination alone does not. Consider a line on which moving one step to the right is a local improvement at every strategy below a . Suppose also that moving one step to the left is a local improvement at every strategy above $a + 1$. Every run ends at a or at $a + 1$. If a pays more in some contingencies and $a + 1$ in others, neither is dominant. The mechanism is then not local improvement solvable under these one-step neighbors, because its two locally unimprovable strategies are not payoff-equivalent. Payoff-equivalent locally unimprovable strategies, on the other hand, are allowed, and they occur in the applications. In the sealed-bid second-price auction with a value on the bid grid, the bid equal to the value and the bid one step above it give the same utility in every contingency. Both bids are locally unimprovable, and both are dominant. In an extensive-form mechanism, two strategies that specify the same action at every situation that both of them reach are payoff-equivalent. The definition therefore allows locally unimprovable strategies that differ only at situations they do not reach, without any further convention.

C.2 Other conditions on a step

Most other worst-best comparison concepts lie between the four conditions of Table 1. To place them, fix a situation h that s_i^* reaches and an action $a' \neq s_h^*$. As in Section 5.1, a continuation after an action at h is the player's actions at the situations that follow the action. Dominance and obvious dominance compare every continuation after a' with the continuation of s_i^* after s_h^* , while action dominance and strong obvious dominance require the comparison for every continuation after s_h^* as well. For action dominance this holds because, by perfect recall, no situation follows both actions, so the common strategy r in (AD) does not link the two continuations. Two families of concepts fill the left column of Table 1. Zhang and Levin

(2018) partition the contingencies and compare the worst case against the best case within each element of the partition. The finest partition gives dominance and the coarsest gives obvious dominance. Arribillaga, Massó and Neme (2024) partition the players, hold fixed the strategies of the players in one's own group, and compare the worst case against the best case over the other players' strategies. The concepts of Pycia and Troyan (2023) fill the bottom row, from strong obvious dominance, where the player plans no future move, to obvious dominance, where she plans all of them.

The conditions of Catonini and Xue (2026) compare two actions at h and treat the player's own continuation in a third way. Their baseline condition fixes a partition of $S(h)_{-i}$. The recommended action dominates a' if every continuation after a' is matched by some continuation after s_h^* . In each element of the partition, the worst payoff of the match must be at least the best payoff of the continuation it matches. With the singleton partition the baseline condition is their weak dominance between actions, and their obvious dominance between actions uses the trivial partition. Because their condition asks only that some match exist, each of the two concepts is implied at h by both conditions in its row of Table 1. The continuation of s_i^* supplies the match for dominance and obvious dominance, and any continuation supplies it for action dominance and strong obvious dominance.

When the two concepts are required at every situation that a strategy reaches, the freedom to choose the match adds nothing to dominance and obvious dominance. Catonini and Xue (2026) show that a strategy is dominant if and only if every action it prescribes at a situation it reaches is weakly dominant, and the same holds for obvious dominance. Since a finer partition gives a weaker condition, their family lies in the left column, between strategy-proofness and OSP.³¹ Action-proofness implies weak dominance between actions at every situation that the recommended strategy reaches, and SOSP implies obvious dominance between actions there, but by Proposition 5(b) neither converse holds.

Their local dominance lies outside the table, since it treats two parts of $S(h)_{-i}$ differently. Where s_h^* ends the player's play, the worst payoff of s_h^* must be at least the best payoff of a' . Elsewhere every continuation after a' must be repeated stage by stage after s_h^* with the same outcome. By their results, local strategy-proofness implies strategy-proofness and coincides with OSP in static mechanisms. In dynamic games, they note, local dominance neither implies nor is implied by obvious or strong obvious dominance. Local strategy-proofness is also independent of action-proofness. The sealed-bid second-price auction is AP and not OSP, so it

³¹The scenario-by-scenario dominance of Catonini and Xue (2026) lets the match change from one element of the partition to another, and their wishful dominance is the singleton case of scenario-by-scenario dominance. They also show that, required at every situation that a strategy reaches, these conditions are equivalent to the baseline ones and so stay in the left column.

is not locally strategy-proof. The English auction with alternating bids, which they argue is locally strategy-proof, is not AP. Consider a type whose value lies above the price she would pay if she bid at her first turn and the rival then left. An action dominant strategy is dominant by Proposition 5(a), so it bids at that turn. Bidding there is worse than leaving at once if she keeps bidding above her value and the rival leaves only after the price passes it. No strategy of such a type is therefore action dominant.

C.3 The ascending clock auction under one-situation neighbors

Under one-situation neighbors the ascending clock auction is not local improvement solvable. As in the two-move game of Figure 3, some strategy can be improved only by changing two actions at once. A strategy specifies an action at every price at which the bidder could still be active, including prices she reaches only if she does not exit when she planned to. Take a value strictly between 0 and $K - 1$ that is not a grid point. Consider a strategy that exits at a price b below the recommended exit price a^* but, if still in, would exit only at a price b' above a^* . This strategy is locally unimprovable. Staying at b alone moves the exit price to b' , which is not an improvement, since she then wins at a price above her value whenever the last rival exits at a price in $[a^*, b')$. Exiting earlier is never better, and a change at a price above b alone has no effect, since the strategy has exited by then. An improvement must change two actions at once, staying at b and exiting at a price above b and no higher than a^* . Under one-step neighbors on exit prices, by contrast, the exit price b has the neighbor $b + 1$, and Proposition 2 applies.

C.4 A proposed design: the multi-account clock

The Ausubel auction asks a bidder at each price for a number of units. A demand stated at a price answers in one number the comparisons of Proposition 3 for her last unit and, if she gives it up, for the units before it. A design for future work would ask these comparisons one at a time, as the ascending clock auction does for a single unit.³²

The multi-account clock is defined as follows. Each bidder has one account for each unit she may receive, L in all, ordered as her marginal values are, from the first unit to the last. All accounts are open when the auction starts. The price rises one grid step at a time, as in the Ausubel auction. At each price the bidder is asked, starting with her last open account, whether she still wants that account's unit at the current price. An answer of no closes the account and moves the question to the account before it, and an answer of yes ends the questions at that

³²This subsection is preliminary. It may be incorporated into a new version of the paper with the results of an experiment that tests this design, which I plan to run in the near future.

price. An account can therefore be closed only after every account after it has been closed, and the bidder's demand at each price is the number of her open accounts. Clinching, payments, and the end of the auction are as in the Ausubel auction, and a bidder observes only the current price and whether the auction is still open.

The multi-account clock and the Ausubel auction implement each other's outcomes. The answers at each price determine the demand there, and every sequence of demands of at most L units that the Ausubel auction allows can be produced by such answers. Maps in both directions therefore send each strategy of either auction to a strategy of the other with the same demand at every price, and sincere bidding to sincere bidding. Sincere bidding, which keeps account ℓ open exactly while the price is below θ^ℓ , is weakly dominant in the multi-account clock, as it is in the Ausubel auction with no bid information. Since a bidder observes only the price and whether the auction is still open, a strategy in either auction amounts to a vector of exit prices, one for each unit. A unit she keeps at every price below K has exit price K . The payoff of the strategy is that of the same vector of marginal bids in the static Vickrey auction. Under one-step neighbors on these vectors, both auctions are therefore local improvement solvable by Proposition 3, for the same reason.

Neither auction is OSP for a bidder who demands two or more units. Reducing her demand by one unit, or closing an account other than the first, at a price below that unit's marginal value keeps her other units in play. The best outcome of that deviation, winning those units at a low price if the rivals' demand falls to zero, exceeds the worst outcome of the recommended action, under which she clinches nothing more. Obvious strategy-proofness on its own therefore predicts no difference between the two auctions.

The two auctions differ in the salience of the local comparison. Each question of the multi-account clock asks the bidder to compare keeping one unit for one more step with giving it up now, the rest of her strategy held fixed. The comparison involves one case, the contingencies in which the unit would clinch at the current price, and the unit's marginal value determines its sign. The Ausubel auction asks for the same comparisons through a single number, so a bidder who states a demand must separate them herself. Since demands that differ by two units are not neighbors, a question of the Ausubel auction is not confined to a choice between two neighboring strategies. The prediction for this design is therefore that the multi-account clock produces more dominant-strategy play than the Ausubel auction, which in turn produces more than the static Vickrey auction.

The multi-account clock is close to the way the Ausubel auction has been run in the laboratory.³³ The laboratory interfaces let a bidder drop units with buttons or keystrokes instead of stating her demand as a number. With two units they differ from the multi-account clock mainly in not asking an explicit question for each open account. As far as I know, no experiment has asked a question for each open account or compared a unit-by-unit interface with one in which the bidder states her demand as a number. A test of the prediction with obvious strategy-proofness held fixed would compare the multi-account clock, the Ausubel auction with stated demands, and the static Vickrey auction. Every bidder would demand two or more units, and no one would see rivals' dropouts, which also excludes the information whose effect Section 7.1 attributes to learning from the environment.

C.5 The pick-an-object serial dictatorship

For the pick-an-object serial dictatorship the recommended strategy is dominant. [Bó and Hakimov \(2024\)](#) state only that it is an equilibrium strategy there. They show, however, that every non-bossy rule with an OSP implementation, serial dictatorship among them, has a pick-an-object implementation in weakly dominant strategies. All agents first pick from the menu of all objects. The agents are then processed in the priority order. Each is assigned her current pick, and every agent with lower priority who picked the same object picks again from the objects that remain. The objects that remain when an agent is processed depend only on the agents with higher priority, who never observe her picks. Picking the most preferred object on each menu therefore gives her the best remaining object whatever the others do. The pick-an-object serial dictatorship therefore implements the outcomes of serial dictatorship with list submission.

³³In the Ausubel auction of [Engelmann and Grimm \(2009\)](#), a bidder drops one unit by clicking one button and both units by clicking another. In the instructions used by [Kagel and Levin \(2009\)](#), each keystroke drops the bidder's lowest-valued remaining unit.

References

- Abdulkadiroğlu, Atila, and Tayfun Sönmez.** 2003. “School Choice: A Mechanism Design Approach.” *American Economic Review*, 93(3): 729–747.
- Archer, Aaron, and Robert Kleinberg.** 2014. “Truthful Germs Are Contagious: A Local-to-Global Characterization of Truthfulness.” *Games and Economic Behavior*, 86: 340–366.
- Arribillaga, R. Pablo, Jordi Massó, and Alejandro Neme.** 2024. “Obvious Strategy-Proofness with Respect to a Partition.” Working paper, arXiv:2409.05315.
- Ashlagi, Itai, and Yannai A. Gonczarowski.** 2018. “Stable Matching Mechanisms Are Not Obviously Strategy-Proof.” *Journal of Economic Theory*, 177: 405–425.
- Ausubel, Lawrence M.** 2004. “An Efficient Ascending-Bid Auction for Multiple Objects.” *American Economic Review*, 94(5): 1452–1475.
- Babaioff, Moshe, Shahar Dobzinski, and Ron Kupfer.** 2023. “Simplicity in Auctions Revisited: The Primitive Complexity.” In *Proceedings of the 24th ACM Conference on Economics and Computation*. 153–182. ACM.
- Banovetz, James, and Ryan Oprea.** 2023. “Complexity and Procedural Choice.” *American Economic Journal: Microeconomics*, 15(2): 384–413.
- Bó, Inácio, and Rustamdjan Hakimov.** 2020. “Iterative versus Standard Deferred Acceptance: Experimental Evidence.” *Economic Journal*, 130(626): 356–392.
- Bó, Inácio, and Rustamdjan Hakimov.** 2024. “Pick-an-Object Mechanisms.” *Management Science*, 70(7): 4693–4721.
- Börger, Tilman, and Jiangtao Li.** 2019. “Strategically Simple Mechanisms.” *Econometrica*, 87(6): 2003–2035.
- Breitmoser, Yves, and Sebastian Schweighofer-Kodritsch.** 2022. “Obviousness around the Clock.” *Experimental Economics*, 25: 483–513.
- Brown, Alexander L., Daniel G. Stephenson, and Rodrigo A. Velez.** 2026. “Testing the Simplicity of Strategy-Proof Mechanisms.” *Economic Theory*, 81(1): 13–42.
- Carroll, Gabriel.** 2012. “When Are Local Incentive Constraints Sufficient?” *Econometrica*, 80(2): 661–686.

- Cason, Timothy N., and Charles R. Plott.** 2014. “Misconceptions and Game Form Recognition: Challenges to Theories of Revealed Preference and Framing.” *Journal of Political Economy*, 122(6): 1235–1270.
- Cason, Timothy N., Tatsuyoshi Saijo, Tomas Sjöström, and Takehiko Yamato.** 2006. “Secure Implementation Experiments: Do Strategy-Proof Mechanisms Really Work?” *Games and Economic Behavior*, 57: 206–235.
- Catonini, Emiliano, and Jingyi Xue.** 2026. “Local Dominance.” Working paper, version of July 2, 2026. Earlier versions at arXiv:2012.14432.
- Chen, Yan, and Tayfun Sönmez.** 2006. “School Choice: An Experimental Study.” *Journal of Economic Theory*, 127: 202–231.
- Chew, Soo Hong, and Wenqian Wang.** 2022. “Generalizing Obvious Dominance Using the Sure-Thing Principle.” Working paper, National University of Singapore.
- Cho, Wonki Jo.** 2016. “Incentive Properties for Ordinal Mechanisms.” *Games and Economic Behavior*, 95: 168–177.
- Dubins, Lester E., and David A. Freedman.** 1981. “Machiavelli and the Gale-Shapley Algorithm.” *American Mathematical Monthly*, 88(7): 485–494.
- Engelmann, Dirk, and Veronika Grimm.** 2009. “Bidding Behaviour in Multi-Unit Auctions: An Experimental Investigation.” *Economic Journal*, 119(537): 855–882.
- Enke, Benjamin, and Cassidy Shubatt.** 2023. “Quantifying Lottery Choice Complexity.” NBER Working Paper 31677.
- Gale, David, and Lloyd S. Shapley.** 1962. “College Admissions and the Stability of Marriage.” *American Mathematical Monthly*, 69(1): 9–15.
- Georganas, Sotiris, Dan Levin, and Peter McGee.** 2017. “Optimistic Irrationality and Overbidding in Private Value Auctions.” *Experimental Economics*, 20: 772–792.
- Glazer, Jacob, and Ariel Rubinstein.** 1996. “An Extensive Game as a Guide for Solving a Normal Game.” *Journal of Economic Theory*, 70: 32–42.
- Gonczarowski, Yannai A., Ori Heffetz, and Clayton Thomas.** 2023. “Strategyproofness-Exposing Mechanism Descriptions.” NBER Working Paper 31506, version of July 2023. arXiv:2209.13148v2. One-page abstract in EC’23.

- Hakimov, Rustamdjan, and Dorothea Kübler.** 2021. “Experiments on Centralized School Choice and College Admissions: A Survey.” *Experimental Economics*, 24: 434–488.
- Harstad, Ronald M.** 2000. “Dominant Strategy Adoption and Bidders’ Experience with Pricing Rules.” *Experimental Economics*, 3: 261–280.
- Hassidim, Avinatan, Assaf Romm, and Ran I. Shorrer.** 2021. “The Limits of Incentives in Economic Matching Procedures.” *Management Science*, 67(2): 951–963.
- Hendon, Ebbe, Hans Jørgen Jacobsen, and Birgitte Sloth.** 1996. “The One-Shot-Deviation Principle for Sequential Rationality.” *Games and Economic Behavior*, 12(2): 274–282.
- Kagel, John H., and Dan Levin.** 2001. “Behavior in Multi-Unit Demand Auctions: Experiments with Uniform Price and Dynamic Vickrey Auctions.” *Econometrica*, 69(2): 413–454.
- Kagel, John H., and Dan Levin.** 2009. “Implementing Efficient Multi-Object Auction Institutions: An Experimental Study of the Performance of Boundedly Rational Agents.” *Games and Economic Behavior*, 66: 221–237.
- 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.
- Kumar, Ujjwal, and Souvik Roy.** 2024. “Local Incentive Compatibility on Gross Substitutes and Other Non-Convex Type-Spaces.” *Journal of Mathematical Economics*, 112: 102989.
- Kumar, Ujjwal, Souvik Roy, Arunava Sen, Sonal Yadav, and Huaxia Zeng.** 2021. “Local Global Equivalence in Voting Models: A Characterization and Applications.” *Theoretical Economics*, 16(4): 1195–1220.
- Li, Shengwu.** 2017. “Obviously Strategy-Proof Mechanisms.” *American Economic Review*, 107(11): 3257–3287.
- Li, Shengwu.** 2024. “Designing Simple Mechanisms.” *Journal of Economic Perspectives*, 38(4): 175–192.
- McGee, Peter, and Dan Levin.** 2019. “How Obvious Is the Dominant Strategy in an English Auction? Experimental Evidence.” *Journal of Economic Behavior and Organization*, 159: 355–365.
- Mennle, Timo, and Sven Seuken.** 2021. “Partial Strategyproofness: Relaxing Strategyproofness for the Random Assignment Problem.” *Journal of Economic Theory*, 191: 105144.

- Mishra, Debasis, Anup Pramanik, and Souvik Roy.** 2016. “Local Incentive Compatibility with Transfers.” *Games and Economic Behavior*, 100: 149–165.
- Nagel, Lea, and Roberto Saitto.** forthcoming. “A Measure of Complexity for Strategy-Proof Mechanisms.” *Econometrica*. Extended abstract in EC’23.
- Niederle, Muriel, and Emanuel Vespa.** 2023. “Cognitive Limitations: Failures of Contingent Thinking.” *Annual Review of Economics*, 15: 307–328.
- Oprea, Ryan.** 2020. “What Makes a Rule Complex?” *American Economic Review*, 110(12): 3913–3951.
- Oprea, Ryan.** 2024. “Decisions under Risk Are Decisions under Complexity.” *American Economic Review*, 114(12): 3789–3811.
- Osborne, Martin J., and Ariel Rubinstein.** 1994. *A Course in Game Theory*. Cambridge, MA: MIT Press.
- Pycia, Marek, and Peter Troyan.** 2023. “A Theory of Simplicity in Games and Mechanism Design.” *Econometrica*, 91(4): 1495–1526.
- Rees-Jones, Alex.** 2018. “Suboptimal Behavior in Strategy-Proof Mechanisms: Evidence from the Residency Match.” *Games and Economic Behavior*, 108: 317–330.
- Roth, Alvin E.** 1982a. “The Economics of Matching: Stability and Incentives.” *Mathematics of Operations Research*, 7(4): 617–628.
- Roth, Alvin E.** 1982b. “Incentive Compatibility in a Market with Indivisible Goods.” *Economics Letters*, 9(2): 127–132.
- Sato, Shin.** 2013. “A Sufficient Condition for the Equivalence of Strategy-Proofness and Nonmanipulability by Preferences Adjacent to the Sincere One.” *Journal of Economic Theory*, 148(1): 259–278.
- Schotter, Andrew, Keith Weigelt, and Charles Wilson.** 1994. “A Laboratory Investigation of Multiperson Rationality and Presentation Effects.” *Games and Economic Behavior*, 6: 445–468.
- Selten, Reinhard.** 1975. “Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games.” *International Journal of Game Theory*, 4(1): 25–55.

- Shapley, Lloyd, and Herbert Scarf.** 1974. “On Cores and Indivisibility.” *Journal of Mathematical Economics*, 1(1): 23–37.
- Troyan, Peter, and Thayer Morrill.** 2020. “Obvious Manipulations.” *Journal of Economic Theory*, 185: 104970.
- Velez, Rodrigo A., and Alexander L. Brown.** 2022. “Empirical Strategy-Proofness.” Working paper, Texas A&M University. Earlier version at arXiv:1907.12408.
- Watson, Joel, and Linchen Zhang.** 2026. “Conditional Dominance and Rationalizability with Information-Based Continuation Reasoning.” Working paper, University of California San Diego.
- Yang, Frank, and Piotr Dworczak.** 2026. “Searchable Menus.” Working paper, arXiv:2608.17446.
- Zhang, Luyao, and Dan Levin.** 2017. “Bounded Rationality and Robust Mechanism Design: An Axiomatic Approach.” *American Economic Review: Papers and Proceedings*, 107(5): 235–239.
- Zhang, Luyao, and Dan Levin.** 2018. “Partition Obvious Preference and Mistrust in Mechanism Design: Theory and Experiment.” Working paper, SSRN 2927190.