Participation Security and Elementary Mechanisms*

Benjamin Brooks Songzi Du

June 3, 2026

Abstract

A mechanism is *participation secure* if each agent has an action that guarantees them a payoff better than their outside option, no matter how other agents behave and no matter how uncertainty about preferences is resolved. Participation security implies that every agent's interim payoff exceeds their outside option, in every information structure and equilibrium. The converse is true if the mechanism is *elementary*: Every action is the unique best response for some type in some information structure and equilibrium. Elementary mechanisms are shown to be a sufficient class for maximizing the welfare *guarantee*. Hence, for maximizing the guarantee, participation security is equivalent to interim participation in all information structures and equilibria.

Keywords: Mechanism design, information design, dual reduction, max-min, Bayes correlated equilibrium, robustness, participation security, individual rationality.

JEL Classification: C72, D44, D82, D83.

*Brooks: Department of Economics, University of Chicago, babrooks@uchicago.edu; Du: Department of Economics, University of California San Diego, sodu@ucsd.edu. We are very grateful for many insightful conversations with Roger Myerson. We also thank Aden Chen for a stimulating discussion that lead to the results in this paper. This research has been supported by NSF #2215259 and #2215475.

1 Introduction

In all of the standard auction formats, e.g., first-price, second-price, and all-pay auctions, each participant has the option to bid zero. The zero bid is exceptional: no matter how others bid and no matter how other uncertainty about preferences is resolved, bidding zero secures a payoff that is at least what a bidder could obtain by not participating in the auction. As a result, no matter what kind of information the bidders have and no matter what equilibrium is being played, all agents are weakly willing to participate in the mechanism, simply because they have the option to participate and bid zero.

This article studies a general purpose model of Bayesian mechanism design, where payoff uncertainty is parametrized by a state of the world. The value of the (possibly state-dependent) outside option is normalized to zero. For a given mechanism, we say that an agent's action is *secure* if it guarantees them a non-negative payoff, no matter what others do and no matter what is the state. The mechanism is *participation secure* if each agent has a secure action. The standard auctions referenced above are all participation secure.

A perhaps more familiar way of modeling participation in Bayesian mechanism design is *(interim) individual rationality*: In equilibrium, each agent's interim expected payoff is non-negative. Participation security implies that no matter what is the information structure and equilibrium, interim individual rationality is satisfied. We refer to this property of a mechanism as *strong individual rationality.* The converse statement is obviously false: Even if interim payoffs are non-negative in all information structures and equilibria, the mechanism may not be participation secure. For example, suppose Ann has an action that unilaterally triggers a doomsday device, which kills Ann and Bob. If Bob leaves the mechanism, he prevents Ann from accessing the device and ensures his own survival. The mechanism is thus not participation secure. But the suicidal action is strictly dominated for Ann and will never be played in equilibrium. It could therefore still be the case that Bob's interim payoff is non-negative in all information structures and equilibria.

This example suggests that while strong individual rationality is more permissive than participation security, it might be so in a superficial way. After all, what is the point of building a mechanism with strictly dominated actions that no one will ever play? One might think that given any strongly individually rational mechanism, it is possible to simplify the mechanism by "trimming away" inessential features, such as strictly dominated actions, and thereby narrow the gap between strong individual rationality and participation security.

In this paper, we show that this intuition is exactly correct for the purposes of *guarantee maximization.* A mechanism's *guarantee* is defined to be minimum expected welfare of the designer across all information structures and equilibria. A recent line of research has studied the maximum guarantee that can be achieved among participation secure mechanisms.[1] This approach to mechanism design avoids the problem of overfitting the mechanism to a particular information structure. As shown in Brooks and Du (2024a), solutions to this problem have a great deal of structure: the guarantee can always be maximized by mechanisms in which actions are linearly ordered, the "lowest" action is secure, and the only relevant equilibrium constraints are those associated with "local" deviations away from the secure action. This exceptional structure has facilitated closed form solutions for settings such as optimal auctions (proportional auctions) and public goods (proportional cost sharing mechanisms).[2]

Could the designer achieve even higher guarantees if participation security is relaxed to strong individually rationality? Our Theorem 1 says that answer is definitively no: Any welfare guarantee that can be achieved with a mechanism that is strongly individually rational can also be achieved with a mechanism that is participation secure.

The proof of the theorem operationalizes the idea of simplifying the mechanism, as we now explain. In general, participation security is more demanding than strong individual rationality, but our Theorem 2 shows that they are equivalent for mechanisms that are *elementary*, meaning that for every agent and action, there is an information structure and equilibrium and some type for which that action is the unique best response. The first-price auction is an example of an elementary mechanism: In the independent private value information structure of Vickrey (1961), each bidder strictly prefers their equilibrium bid to all other bids, and every bid could be an equilibrium action for an appropriately chosen value distribution.

Elementary mechanisms are a natural and immediate generalization of the notion of elementary games of Myerson (1997). By adapting Myerson's analysis, we show that starting from any mechanism, it is possible to obtain from it an elementary mechanism via a process called iterated dual reduction. This process generates a sequence of mechanisms, where each is obtained from its predecessor by partitioning each agent's actions in a certain way, and replacing the actions in a given cell with a single carefully chosen mixed action.

[1]See, e.g., Du (2018); Bergemann, Brooks, and Morris (2016); Brooks and Du (2021, 2023, 2024a,b); Brooks, Du, and Haberman (2024); Brooks, Du, and Zhang (2026b); Brooks, Du, and Feffer (2026a). Brooks and Du (2025) provides an overview of this research program and the justification for (and critique of) this approach to mechanism design.

[2]This analytical simplification depends critically on the use of participation security: It is straightforward to check whether the first action in the sequence is secure. By contrast, strong individual rationality involves universal quantification over all information structures and equilibria.

A crucial property of dual reduction is that for every information structure, it shrinks the set of equilibrium outcomes: for every equilibrium of the reduced mechanism, there is a corresponding equilibrium of the original mechanism which induces the same interim payoffs for the agents and the same joint distribution over states and outcomes. As a result, dual reduction preserves strong individual rationality, and it also raises the guarantee. A fortiori, the same is true of an elementary mechanism that is obtained as a limit of dual reductions. Thus, in maximizing the guarantee subject to strong individual rationality, it is without loss to restrict attention to elementary mechanisms.

Moreover, in an elementary mechanism, all actions can be made to be strict best responses (and therefore are played with strictly positive probability). This implies that every interim belief about the state and others' actions can be realized as an interim belief in some information structure and equilibrium. That fact, combined with strong individual rationality and the minimax theorem, delivers the conclusion that each agent must have a secure action.

We have emphasized the equivalence of participation security and strong individual rationality for guarantee maximization. But as just explained, the equivalence holds whenever the designer uses elementary mechanisms. If a mechanism is not elementary, then there must be some way of deviating for some agent, such that the agent weakly prefers to deviate no matter the information structure and equilibrium. Statements about equilibrium that depend on whether the agents would follow such persistent deviations seem especially dubious. Dual reduction eliminates such forms of fragility in equilibrium behavior. Note that there can still be equilibria of elementary mechanisms in which an agent has more than one best response. But in a certain sense, iterated dual reduction eliminates as much of this kind of indeterminacy as possible, and they are desirable whenever the designer would, all else equal, prefer mechanisms with smaller sets of equilibria.[3]

[3]At first blush, it seems that certain mechanisms that feature prominently in the literature on implementation under complete information are not elementary, e.g., those of Maskin (1999) and Abreu and Matsushima (1992). Indeed, in all equilibria, the agents only use particular action profiles in which they "agree" in their reports about the state. However, this statement only concerns behavior in equilibria in which the agents have complete information about the state. For example, in Abreu and Matsushima (1992), each agent's action consists of many distinct "reports" of the state. For a given "baseline" report, the agent is incentivized to report what the state is according to their own first order beliefs, but in other reports, they are incentivized to agree with others' reports. Under complete information, the unique rationalizable action profile involves all reports of all agents being equal to the true state. But analyzing the same mechanism under incomplete information, if Ann assigned probability close to 1 that the state is θ and that Bob assigns probability close to 1 to θ', then it could be a best response for Ann to make a baseline report of θ but make a different report equal to θ' to match Bob's anticipated baseline report. Indeed, we suspect that by appropriately constructing higher order beliefs, it may be possible to show that Abreu-Matsushima mechanisms are in fact elementary.

Stepping even further back, participation constraints represent a specific form of moral hazard, wherein exactly one agent unilaterally decides to "leave" the mechanism. We do not model explicitly all of the actions available to agents outside of the mechanism. It may very well be that these unilateral exits are just one aspect of a wider game in which agents affect one another through choices made outside the mechanism, and the payoff from unilateral exit may be different from what happens if more than one agent refuses to participate. For example, in bilateral trade, just because agents choose not to trade through the designer's platform, that does not mean that they cannot trade through some other rival platform. When we reduce all of those features of the environment to an interim participation constraint, we are implicitly selecting for equilibria in which the agents coordinate on participating in the designer's mechanism. This is also true with participation security.

In Section 6, we discuss the possibility that the "outside option" represents a specific outcome that is common to all agents, like that the mechanism simply ceases to exist, and that this outcome is something that the designer could themself implement. In that case, one could build secure mechanisms in which the secure action triggers this outcome, say, by shutting down the mechanism altogether. But that need not be the case; in the exchange context, it could be that the secure action does allow for trade to occur, but just on terms that are guaranteed to be favorable to the agent playing the secure action (as in Brooks and Du, 2024b). Even though in the broader game there might be no-trade equilibria in which all agents refuse to trade, by excluding a categorical "no-trade" action from the mechanism itself, we are implicitly selecting for equilibria in which the agents permit trade to occur in certain circumstances. Participation security is therefore a flexible modeling tool, that allows us to select for equilibria with full participation, while still acknowledging that participation is a voluntary choice of the agents. By contrast, the standard approach in Bayesian mechanism design is to select a *particular* equilibrium under which agents are willing to participate, even though there might be other equilibria within the mechanism from which agents would have an incentive to exit. In that sense, mechanisms that are participation secure or strongly individually rational are less susceptible to unraveling than mechanisms for which only certain equilibria are interim individually rational.

The preceding paragraphs describe just one approach to robust mechanism design. A distinct approach, going back to Dasgupta, Hammond, and Maskin (1979), models robustness as ex post implementability. This implementation concept relies on further structure on the agents' information and preferences. In particular, it is assumed that each agent has a "payoff type," which they know for sure, and that the vector of payoff types determines all of the agents' and the designer's ex post preferences. In addition to the payoff type, agents might also have additional information about others' payoff types and information. In an auction setting, the payoff type might be a private value for the good being sold. Additionally, the implementation concept allows the designer to select their preferred equilibrium, but the equilibrium and participation constraints have to hold ex post, conditional on the realized vector of all agents' payoff types. Bergemann and Morris (2005) and Chung and Ely (2007) identify conditions under which ex post implementation is without loss,[4] but both papers also give examples of settings in which a designer could achieve more with Bayesian mechanisms than they are able to with ex post implementable mechanisms.[5]

In our model, there are no payoff types; there is only interim information about the payoff state. In fact, one possibility is that the agents do not know anything about the state at all.[6] As a result, there is nothing that the agents know for sure and with respect to which incentives could be provided ex post. Strong individual rationality and participation security are therefore useful concepts when the designer does not even know the possible payoff types that agents might have. Note that neither strong individual rationality or participation security imply that equilibrium outcomes under any particular information structure are ex post individually rational.

Haberman and Jagadeesan (2025) study optimal auctions with correlated private values. They analyze what can be achieved with Bayesian mechanisms that give the agents "withdrawal rights," meaning, the ability to withdraw from the auction after allocations and payments have been determined. Withdrawal is analogous to a secure action, in that any agent can at any time choose to withdraw and obtain the outside option. Haberman and Jagadeesan (2025) develop and apply a "withdrawal-proofness principle," which selects for equilibria in which the agents never withdraw. By contrast, our analysis considers equilibria in which agents may take their secure actions.

As mentioned above, our analysis builds on the pioneering work of Myerson (1997). Myerson (2024) broadens that program to the analysis of communication equilibrium in

[4]Chung and Ely (2007) analyze a private value auction problem, in which case ex post implementation reduces to dominant strategy implementation and ex post individual rationality. See also Chen and Li (2018) and Yamashita and Zhu (2018) for generalizations of Chung and Ely's analysis to other mechanism design settings.

[5]Other work in this tradition includes Yamashita (2018), Che (2022), He and Li (2022), and He et al. (2024).

[6]No information would be the worst case if there were no conflict of interest between the designer and the agents. But in general, there are information structures that are more challenging for the designer than no information. Brooks and Du (2024b) give an example from bilateral trade: it may be that there are always gains from trade ex post, so that under no information, a welfare maximizing designer would want the agents to trade for sure, and that outcome is implementable under no information. But there are other information structures for which no interim individual rational and incentive compatible mechanism implements trade with positive probability.

Bayesian games. Our application of dual reductions to mechanisms is a technically superficial extension of Myerson's (1997) treatment of complete information games. Myerson describes dual reduction as a methodology for selecting models of strategic interaction that avoid statements about behavior that cannot be generated with strict incentives. By contrast, we view dual reduction as a tool for simplifying mechanisms in a way that yields favorable guarantees, and with transparent incentives to participate.

Brooks and Du (2024a) also use dual reductions to study guarantee maximization subject to participation security. There are some important differences. Brooks and Du (2024a) develop a notion of dual reduction that is distinct from the dual reduction introduced in Myerson (1997) (and that is adapted in this paper). In particular, Brooks and Du (2024a) studies the dual of the linear program of minimizing welfare in a given mechanism across all Bayes correlated equilibria (Bergemann and Morris, 2013, 2016). The dual variables are then interpreted as a "most tempting" stochastic deviation for the agents. The dual reduction is defined to be the restriction of the original mechanism to the path of mixed actions generated by starting from a secure action and iteratively applying the most tempting deviation. Brooks and Du (2024a) show that dual reduction necessarily increases the guarantee. More importantly, it demonstrates that the optimal guarantee can be obtained by considering mechanisms in which each agent's actions are ordered in a sequence, requiring the first action in the sequence to be secure, and by maximizing a lower bound on the guarantee that is implied by local equilibrium constraints. This provides a foundation for the first-order methodology for guarantee maximization that was introduced and applied by Brooks and Du (2021, 2024b) and others referenced below.

A contribution of this paper is to provide a foundation for the use of participation security; all of the mechanisms derived by the papers in this literature are also solutions to the program of maximizing the guarantee among mechanisms that are strongly individually rational. The list of such mechanisms includes proportional auctions (Brooks and Du, 2021), proportional cost-sharing mechanisms (Brooks and Du, 2023), proportional price trading mechanisms (Brooks and Du, 2024b), compound proportional auctions (Brooks, Du, and Feffer, 2026a), and market order mechanisms (Brooks, Du, and Zhang, 2026b).

In an addendum, we explain how dual reduction can also be applied to information design, and specifically the problem of finding the information structures that are most challenging for the designer. In particular, dual reduction of information structures shrinks the set of equilibrium outcomes for all mechanisms. Iterative dual reduction of an information structure leads to one that is *elementary*, meaning that for every type, there is a mechanism and equilibrium in which that type strictly prefers their strategy to that of any other type. An implication of this result is that it is without loss to restrict attention to elementary information structures when minimizing the *potential*: The highest welfare that can be achieved in any participation secure mechanism and equilibrium.

The rest of this paper is structured as follows: Section 2 lays out the model of Bayesian mechanism design. Section 3 states our main result on participation constraints and guarantee maximization. Section 4 describes the process of dual reduction, elementary mechanisms, and their properties. Section 5 presents the result that strong individual rationality is equivalent to participation security for elementary mechanisms, which then is used to prove the main theorem on guarantee maximization. Section 6 is an addendum on applications of dual reduction to information design and welfare potentials.

2 Model

There are agents $i = 1, \ldots, N$ and a designer.

The designer controls an outcome $\omega \in \Omega$. All sets are finite unless otherwise stated.

Agents and the designer have expected utility preferences over outcomes and a *state* $\theta \in \Theta$. Agent i's utility index is $u_i(\omega, \theta)$ and the designer's index is $w(\omega, \theta)$.

Each agent can leave the mechanism. We interpret u_i as their payoff net of the payoff from leaving, which may be state-dependent.

An *information structure* consists of: a set of signals S_i for each agent i; and a joint probability distribution $\sigma(s, \theta)$ over

$$\underbrace{S_1 \times \cdots \times S_N}_{\equiv S} \times \Theta.$$

With a slight abuse of notation, we sometimes adopt a shorthand of denoting an information structure (S, σ) by its distribution σ.

A *mechanism* consists of: a set of actions A_i for each agent i; and a mapping

$$m : \underbrace{A_1 \times \cdots \times A_N}_{\equiv A} \to \Delta(\Omega).$$

We sometimes adopt a shorthand of denoting a mechanism (A, m) by its mapping m.

Given an information structure and mechanism, a *(behavioral) strategy* for agent i is a mapping $\beta_i : S_i \to \Delta(A_i)$. Given a strategy profile β, we write

$$\beta(a|s) = \prod_i \beta_i(a_i|s_i).$$

The expected utilities of agent i and the designer are:

$$U_i(\beta; m, \sigma) = \sum_{s,\theta,a,\omega} u_i(\omega, \theta) m(\omega|a) \beta(a|s) \sigma(s, \theta);$$
$$W(\beta; m, \sigma) = \sum_{s,\theta,a,\omega} w(\omega, \theta) m(\omega|a) \beta(a|s) \sigma(s, \theta).$$

The strategy profile β is a *(Bayes Nash) equilibrium* of the game (m, σ) if for all i and β_i', $U_i(\beta; m, \sigma) \geqslant U_i(\beta_i', \beta_{-i}; m, \sigma)$. Let $E(m, \sigma)$ be the set of equilibria of (m, σ). Since (m, σ) is a finite game, $E(m, \sigma)$ is always non-empty.

3 Participation Constraints and Welfare Guarantees

There is a prior over Θ, denoted by μ. Suppose μ has full support, i.e., $\mu(\theta) > 0$ for all θ. We say that the information structure σ is *compatible with* μ if for all θ, $\sum_s \sigma(s, \theta) = \mu(\theta)$. In other words, the marginal distribution on θ is equal to μ. Let $\mathcal{I}(\mu)$ be the set of all information structures compatible with μ.

Given an equilibrium β of a game (m, σ), for each i and $s_i \in S_i$ let us define

$$U_i(s_i, \beta; m, \sigma) = \sum_{s_{-i},\theta,a,\omega} u_i(\omega, \theta) m(\omega|a) \beta(a|s) \sigma(s, \theta).$$

Up to rescaling by the marginal likelihood of s_i, this is the interim expected payoff of agent i when they have signal s_i. The mechanism (A, m) is *strongly individually rational* with respect to μ if for every information structure $\sigma \in \mathcal{I}(\mu)$, every equilibrium $\beta \in E(m, \sigma)$, every agent i, and every signal s_i, we have $U_i(s_i, \beta; m, \sigma) \geqslant 0$.

Lemma 1. *If a mechanism is strongly individually rational with respect to a full support prior μ, it is strongly individually rational with respect to any other full support prior μ'.*

Proof. Given any two full support priors μ and μ', there exists an $\rho \in (0, 1)$ such that $\mu = \rho\mu' + (1 - \rho)\mu''$, where μ'' is another full support prior. If a mechanism (A, m) is not strongly individually rational with respect to μ', then it is not strongly individually rational with respect to μ: Let $(S', \sigma') \in \mathcal{I}(\mu)$ and $\beta' \in E(m, \sigma')$ be such that $U_i(s_i', \beta'; m, \sigma') < 0$ for some i and s_i', and pick any $(S'', \sigma'') \in \mathcal{I}(\mu'')$ and any $\beta'' \in E(m, \sigma'')$. Let (S, σ) be the public randomization between (S', σ') and (S'', σ''), with probabilities ρ and $1 - \rho$ (so $S_i = S_i' \cup S_i''$), and let β be the strategy profile where β' is played if (S', σ') is realized, and β'' is played if (S'', σ'') is realized. Then we have $(S, \sigma) \in \mathcal{I}(\mu)$, $\beta \in E(m, \sigma)$, and $U_i(s_i', \beta; m, \sigma) = U_i(s_i', \beta'; m, \sigma') < 0$. □

Thus, we say that a mechanism is strongly individually rational if it is strongly individually rational with respect to a full support prior.

In classical Bayesian mechanism design, there is a fixed information structure that is known to the designer, and interim individual rationality is only required in the particular equilibrium selected by the mechanism designer. Such a mechanism might have other equilibria in which agents receive negative interim payoffs and therefore have an incentive to leave the mechanism. In that case, the conclusion that agents would be willing to participate depends on which equilibrium is selected. Moreover, the agents' willingness to participate may also be dependent on the information structure being the particular one hypothesized by the designer. But the designer might not know precisely the information structure and might not be able to control which equilibrium is played. By implementing a mechanism which is strongly individually rational, the designer *guarantees* that agents will be willing to participate, no matter the environment and equilibrium.

As mentioned in the introduction, a sufficient condition for strong individual rationality (and one which eliminates the universal quantification over information structures and equilibria) is the following: Given a mechanism (A, m), a mixed action $x_i \in \Delta(A_i)$ is *secure* if for all θ and a_{-i},

$$\sum_{\omega, a_i} u_i(\omega, \theta) m(\omega | a_i, a_{-i}) x_i(a_i) \geqslant 0,$$

The mechanism (A, m) is *participation secure* if every agent i has a secure mixed action.[7] Clearly, every participation secure mechanism is strongly individually rational, but not vice versa.

The designer is ultimately concerned not just with participation but also with how much welfare is generated by the mechanism. And the consideration of what might happen in all mechanisms and information structures natural leads to the notion of a mechanism's *welfare guarantee*:

$$G(m) = \inf_{(S,\sigma) \in \mathcal{I}(\mu)} \inf_{\beta \in E(m,\sigma)} W(\beta; m, \sigma).$$

Our main result is that

[7]The literature has typically defined a mechanism to be participation secure if every agent i has a secure pure action. This difference is immaterial for guarantee maximization because we can always enlarge the action space of a mechanism to convert a mixed action into a pure action without changing the set of equilibrium outcomes under any information structure.

Theorem 1. *Let $\mathcal{M}$ be the set of strongly individually rational mechanisms, and let $\mathcal{M}_0$ be the set of participation secure mechanisms. Then*

$$\sup_{(A,m)\in\mathcal{M}} G(m) = \sup_{(A,m)\in\mathcal{M}_0} G(m).$$

4 Elementary Mechanisms

To prove Theorem 1, we adapt the concepts of dual reduction and elementary games introduced by Myerson (1997).

Given a mechanism (A, m), a *direct information structure* is one in which $S_i = A_i$ for all i. Given such an information structure, the truthful strategies are the ones for which $\beta_i(a_i|a_i) = 1$ for all i and a_i. A *Bayes correlated equilibrium* (BCE) of (A, m) is a direct information structure with $S_i = A_i$ for all i and such that the truthful strategies are an equilibrium, meaning that for all i, a_i, and a_i',

$$\sum_{a_{-i},\theta,\omega} u_i(\omega,\theta)[m(\omega|a_i,a_{-i}) - m(\omega|a_i',a_{-i})]\sigma(a_i,a_{-i},\theta) \geqslant 0. \tag{1}$$

Note that in contrast to the definition of BCE given in Bergemann and Morris (2013, 2016), we do not fix a marginal distribution on θ. Rather, the distribution of states is allowed to "float" as in Brooks, Du, and Haberman (2024).

A BCE has *elementary incentives* if (1) is strict for all for every i, $a_i \neq a_i'$. (Such a BCE necessarily has full support on the actions.) The mechanism (A, m) is *elementary* if it has a BCE with elementary incentives.

We now describe a procedure, adapted from Myerson (1997), by which any mechanism can be *dually reduced* to an elementary mechanism.

Consider the following linear programming problem:

$$\begin{aligned}
&\max_{\sigma(a,\theta)\geqslant 0,\nu_i(a_i)\in\mathbb{R}} \sum_i \sum_{a_i} \nu_i(a_i) \\
\text{s.t. } &\sum_a \sigma(a,\theta) = 1 && [\lambda] \\
&\sum_{a_{-i},\theta,\omega} u_i(\omega,\theta)[m(\omega|a_i,a_{-i}) - m(\omega|a_i',a_{-i})]\sigma(a_i,a_{-i},\theta) \geqslant \nu_i(a_i)\ \forall i, a_i, a_i' && [\alpha_i(a_i'|a_i)].
\end{aligned}$$

The variables in square brackets are the Lagrange multipliers on the respective constraints.

We regard the preceding linear program as being primal. The corresponding dual program is

$$
\begin{aligned}
&\min_{\alpha_i(a_i'|a_i)\geqslant 0, \lambda\in\mathbb{R}} \lambda \\
\text{s.t. } &\sum_{i,a_i',\omega} \alpha_i(a_i'|a_i) u_i(\omega,\theta) \left[m(\omega|a_i', a_{-i}) - m(\omega|a) \right] + \lambda \geqslant 0 \ \forall a, \theta \qquad [\sigma(a,\theta)] \\
&\sum_{a_i'} \alpha_i(a_i'|a_i) = 1 \ \forall i, a_i \qquad [\nu_i(a_i)].
\end{aligned}
$$

An α that is feasible for the dual program may be interpreted as a collection of transition probability matrices on actions, one for each agent i.

Because the set of BCE is non-empty and compact, the primal program is feasible and bounded, and hence there exists a saddle point $(\sigma^*, \nu^*, \alpha^*, \lambda^*)$. Moreover, because all obedience constraints (1) are satisfied in a Bayes correlated equilibrium, the optimal value must be equal to zero, and hence $\nu_i^*(a_i) = 0$ and $\lambda^* = 0$. Thus, the dual solution satisfies for all θ and a,

$$
\sum_{i,a_i',\omega} \alpha_i^*(a_i'|a_i) u_i(\omega,\theta) \left[m(\omega|a_i', a_{-i}) - m(\omega|a) \right] \geqslant 0 \tag{2}
$$

A mixed action for agent i is denoted by $x_i \in \Delta(A_i)$. Given an α that is dual feasible (i.e., a collection of transition probability matrices), the mixture is α_i-*stationary* if for all a_i,

$$
x_i(a_i) = \sum_{a_i'} \alpha_i(a_i|a_i') x_i(a_i').
$$

Clearly, the set of α_i-stationary mixtures is a convex polytope.

The α-reduction of the mechanism (A, m) is the mechanism $(\widehat{A}, \widehat{m})$, for which $\widehat{A}_i$ is the set of extreme points of the α_i-stationary mixtures, and

$$
\widehat{m}(\omega|x) = \sum_a m(\omega|a) \prod_i x_i(a_i).
$$

In other words, the α-reduction of (A, m) is the mechanism in which agents are restricted to playing α_i-stationary mixtures.

The next proposition says that for any α that is part of a saddle point, under any information structure, any equilibrium of the α-reduction lifts to an equilibrium of the original game.

Proposition 1. *Let α^* be part of a solution to the dual program for (A, m), and let $(\widehat{A}, \widehat{m})$ be the α^*-reduction. For any information structure (S, σ) and equilibrium $\widehat{\beta}$ of $(\widehat{m}, \sigma)$, the following strategies are an equilibrium of (m, σ):*

$$\beta_i(a_i|s_i) = \sum_{x_i \in \widehat{A}_i} x_i(a_i)\widehat{\beta}_i(x_i|s_i). \tag{3}$$

Moreover, $U_i(s_i, \beta; m, \sigma) = U_i(s_i, \widehat{\beta}; \widehat{m}, \sigma)$ for all i and s_i, and $W(\beta; m, \sigma) = W(\widehat{\beta}; \widehat{m}, \sigma)$.

The proof of the proposition follows closely the proof of Theorem 1 of Myerson (1997):

Proof. We will prove that β is an equilibrium of (m, σ); the rest of the theorem follows by the construction of β and $\widehat{m}$.

For any strategy β'_i for (m, σ), let $\alpha^*_i \circ \beta'_i$ be the strategy defined by

$$(\alpha^*_i \circ \beta'_i)(a_i|s_i) = \sum_{a'_i} \alpha^*_i(a_i|a'_i)\beta'_i(a'_i|s_i).$$

We have $\alpha^*_j \circ \beta_j = \beta_j$ for every agent j. Because of (2), this implies

$$U_i(\alpha^*_i \circ \beta'_i, \beta_{-i}; m, \sigma) \geqslant U_i(\beta'_i, \beta_{-i}; m, \sigma).$$

Thus,

$$U_i\left(\tilde{\alpha}_i \circ \beta'_i, \beta_{-i}; m, \sigma\right) \geqslant U_i(\beta'_i, \beta_{-i}; m, \sigma),$$

where $\tilde{\alpha}_i(a_i, a'_i) = \mathbb{I}_{a_i = a'_i}/2 + \alpha^*_i(a_i, a'_i)/2$.

The Markov chain defined by $\tilde{\alpha}_i$ is aperiodic, and a mixture is $\tilde{\alpha}_i$-stationary if and only if it is α^*_i-stationary. Moreover $(\tilde{\alpha}_i)^n \circ \beta'_i$, as $n \to \infty$, converges to a α^*_i-stationary mixture, where $(\tilde{\alpha}_i)^n$ is the n-th power of $\tilde{\alpha}_i$ (i.e., $(\tilde{\alpha}_i)^n \circ \beta'_i = \tilde{\alpha}_i \circ ((\tilde{\alpha}_i)^{n-1} \circ \beta'_i)$; see for example Stroock (2014), equation (4.1.15). If we take β'_i to be a best response to β_{-i} in (m, σ), this implies that an α^*_i-stationary mixture is also a best response to β_{-i}. Since $\widehat{\beta}_i$ is a best response to $\widehat{\beta}_{-i}$ in $(\widehat{m}, \sigma)$, and β is lifted to m from $\widehat{\beta}$ in $\widehat{m}$, this proves that β_i is a best response to β_{-i} in (m, σ). $\square$

Thus, dual reduction effectively shrinks the set of equilibrium outcomes of the mechanism. Mathematically, this result is an almost immediate consequence of the argument given in Myerson (1997). Myerson's result is that for complete information games, dual reduction effectively shrinks the set of correlated equilibria. In our setting with incomplete information, the analogous statement would be that dual reduction shrinks the set of BCE. However, the proposition shows that Myerson's argument implies an even stronger conclusion: for every information structure, dual reduction shrinks the set of Bayes Nash equilibrium outcomes and the possible equilibrium interim payoff profiles. This observation is important for our subsequent application of dual reduction to guarantee maximization subject to interim participation constraints.

As in Myerson (1997), we can iteratively reduce the mechanism (A, m). The extreme α_i^*-stationary mixtures correspond one-to-one with the recurrent classes of α_i^* (Stroock (2014), Theorem 4.1.10). At each round, as long as $\alpha_i^*(a_i'|a_i) > 0$ for some $a_i' \neq a_i$, either a_i and a_i' are both recurrent states of α_i^*, or a_i is a transient state of α_i^*. In the former case, dual reduction will remove a_i and a_i' from the game (along with all other actions in their recurrence class), and in their place will be an extreme α_i^*-invariant mixture with support equal to their recurrent class. In the latter case, dual reduction will remove the transient a_i without replacement. Thus, every round of non-trivial dual reduction decreases the number of actions of some agent by at least one. This proves that there can be at most finitely many rounds of dual reduction.

Note that there is not a unique path of iterated dual reduction, since at each round, there may be more than one α^* that solves the dual program.

When iterated dual reduction converges to a mechanism $(\widehat{A}, \widehat{m})$, it must be that for that mechanism, the unique optimal dual solution must be $\alpha_i^*(a_i|a_i) = 1$. By strong complementary slackness in linear programming (Goldman and Tucker, 1956, Corollary 2B), there must exist a solution of the primal program for which all obedience constraints are strict, which implies that all actions are played with positive probability. We obtain the following result:

Proposition 2. *For any mechanism (A, m), there exists an elementary mechanism $(\widehat{A}, \widehat{m})$, which is obtained from (A, m) via iterative dual reduction. For any information structure (S, σ) and equilibrium $\widehat{\beta}$ of $(\widehat{m}, \sigma)$, there is an equilibrium β of (m, σ) with $U_i(s_i, \beta; m, \sigma) = U_i(s_i, \widehat{\beta}; \widehat{m}, \sigma)$ for all i and $s_i \in S_i$, and $W(\beta; m, \sigma) = W(\widehat{\beta}; \widehat{m}, \sigma)$.*

This result follows from essentially the same argument as that underlying Theorem 2 of Myerson (1997), adapted to the case with the payoff-state θ. But again, the stated result is stronger: Not only does the reduction to an elementary mechanism shrink the set of (Bayes) correlated equilibria, but it also shrinks the set of equilibrium outcomes and interim utility profiles, for every information structure.

5 Equivalence for Elementary Mechanisms

Our next result reveals why the interest in participation constraints naturally led us to the investigation of elementary mechanisms:

Theorem 2. *An elementary mechanism is strongly individually rational if and only if it is participation secure.*

Proof. We will prove the only if direction, since the if direction is immediate and does not need the hypothesis of an elementary mechanism.

Let (A, m) be an elementary mechanism, and let σ be a BCE in m with elementary incentives. By perturbing σ slightly, we may assume σ has both elementary incentives and a full support on $A \times \Theta$, i.e., $\sigma(a, \theta) > 0$ for every (a, θ).

For the sake of contradiction, suppose m is not participation secure: suppose agent i has no secure mixed action. Consider the zero-sum game played between agent i, who is choosing their action a_i to maximize their payoff, and an adversary who is choosing (a_{-i}, θ) to minimize agent i's payoff. By the minimax theorem, this game has a value, and since agent i does not have a secure action, the value must be negative. As a result, there exists a strategy of the adversary, denoted $\gamma \in \Delta(A_{-i} \times \Theta)$, such that

$$\sum_{\theta, a_{-i}, \omega} u_i(\omega, \theta) m(\omega|a) \gamma(a_{-i}, \theta) < 0$$

for all a_i. Let $\overline{a}_i \in A_i$ be a best response of agent i to the belief γ.

Let $\sigma'(a, \theta) = \gamma(a_{-i}, \theta)$ if $a_i = \overline{a}_i$, and 0 otherwise. We write $\sigma = \epsilon\sigma' + (1 - \epsilon)\sigma''$, where $\epsilon > 0$ is sufficiently small so that $\sigma'' \in \Delta(A \times \Theta)$ and

$$\sum_{a_{-i}, \theta, \omega} u_i(\omega, \theta)[m(\omega|a_i, a_{-i}) - m(\omega|a'_i, a_{-i})]\sigma''(a_i, a_{-i}, \theta) > 0,$$

for every $a_i \neq a'_i$.

Now, consider the following information structure: Nature first chooses σ' with probability ϵ and chooses σ'' with probability $1 - \epsilon$; then nature draws (a, θ) from either σ' or σ'' and informs each agent j their a_j; moreover, agent i is notified whether Nature chose σ' or σ''. By construction, for this information structure each agent j obeying their action recommendation is an equilibrium. In this equilibrium, agent i obtains a negative expected payoff if he is notified of σ' and recommended to play $\overline{a}_i$. This implies that the mechanism m is not strongly individually rational. □

Finally, we use Theorem 2 to obtain the result that the same maximum guarantee is attained with participation security as with strong individual rationality:

Proof of Theorem 1. For any strongly individually rational mechanism (A, m), by Proposition 2 its iterative dual reduction $(\widehat{A}, \widehat{m})$ is elementary, strongly individually rational, and has a weakly higher guarantee. By Theorem 2, $(\widehat{A}, \widehat{m})$ is participation secure. Thus,

$$\sup_{(A,m)\in\mathcal{M}} G(m) \leqslant \sup_{(\widehat{A},\widehat{m})\in\mathcal{M}_0} G(m).$$

The other direction of the inequality follows from $\mathcal{M}_0 \subset \mathcal{M}$. □

An immediate consequence of Theorem 1 and Proposition 2 is that in maximizing the welfare guarantee, it is also without loss to restrict attention to elementary mechanisms:

Corollary 1. *Let $\widehat{\mathcal{M}}$ be the set of elementary mechanisms that are strongly individually rational (equivalently, participation secure). We have*

$$\sup_{(A,m)\in\mathcal{M}} G(m) = \sup_{(A,m)\in\widehat{\mathcal{M}}} G(m) = \sup_{(A,m)\in\mathcal{M}_0} G(m).$$

6 Addendum: Dual Reduction of Information Structures

As a last topic, we describe analogous notions of dual reduction and elementality for information structures. Given an information structure (S, σ), a *direct mechanism* (A, m) is one in which $A_i = S_i$ for all i. Given such a mechanism, the truthful strategies are the ones for which $\beta_i(s_i|s_i) = 1$ for all i and s_i. A direct mechanism (S, m) for (S, σ) is *incentive compatible* if the truthful strategies are an equilibrium, meaning that for all i, s_i, and s_i',

$$\sum_{s_{-i},\theta,\omega} u_i(\omega, \theta)[m(\omega|s_i, s_{-i}) - m(\omega|s_i', s_{-i})]\sigma(s_i, s_{-i}, \theta) \geqslant 0. \tag{4}$$

An incentive compatible direct mechanism (S, m) has *elementary incentives* if (4) is strict for all for every i, $s_i \neq s_i'$. The information structure (S, σ) is *elementary* if it has an incentive compatible direct mechanism with elementary incentives. Examples of elementary information structures include the independent private values auction model of Vickrey (1961) and the affiliated values model of Milgrom and Weber (1982). A trivial example of an information structure that is not elementary is given below.

We now describe a procedure, the analogue of the one in Section 4, by which any information structure can be *dually reduced* to an elementary information structure.

Consider the following linear programming problem:

$$\max_{m(\omega|s)\geqslant 0, \nu_i(s_i)\in\mathbb{R}} \sum_i \sum_{s_i} \nu_i(s_i)$$

$$\text{s.t. } \sum_{\omega} m(\omega|s) = 1 \qquad [\gamma(s)]$$

$$\sum_{s_{-i},\theta,\omega} u_i(\omega,\theta)[m(\omega|s_i,s_{-i}) - m(\omega|s_i',s_{-i})]\sigma(s_i,s_{-i},\theta) \geqslant \nu_i(s_i)\ \forall i, s_i, s_i' \qquad [\alpha_i(s_i'|s_i)].$$

The variables in square brackets are the Lagrange multipliers on the respective constraints.

We regard the preceding linear program as being primal. The corresponding dual program is

$$\min_{\alpha_i(s_i'|s_i)\geqslant 0, \gamma(s)\in\mathbb{R}} \sum_s \gamma(s)$$

$$\text{s.t. } \sum_{i,\theta} u_i(\omega,\theta)\left(\sum_{s_i'} \alpha_i(s_i|s_i')\sigma(s_i',s_{-i},\theta) - \sigma(s,\theta)\right) + \gamma(s) \geqslant 0\ \forall \omega, s \qquad [m(\omega|s)]$$

$$\sum_{s_i'} \alpha_i(s_i'|s_i) = 1\ \forall i, a_i \qquad [\nu_i(s_i)].$$

As in the dual reduction for mechanism, an α that is feasible for the dual program may be interpreted as a collection of transition probability matrices on actions, one for each agent i.

Because the set of direct mechanism is non-empty and compact, the primal program is feasible and bounded, and hence there exists a saddle point $(m^*, \nu^*, \alpha^*, \gamma^*)$. Moreover, because all incentive compatibility constraints (4) can be satisfied in a direct mechanism, the optimal value must be equal to zero, and hence $\nu_i^*(s_i) = 0$ and $\sum_s \gamma^*(s) = 0$. Thus, the dual solution satisfies for all ω and s,

$$\sum_{i,\theta} u_i(\omega,\theta)\left(\sum_{s_i'} \alpha_i^*(s_i|s_i')\sigma(s_i',s_{-i},\theta) - \sigma(s,\theta)\right) + \gamma^*(s) \geqslant 0. \qquad (5)$$

Note that the above equation is still satisfied if $\alpha_i^*(s_i'|s_i)$ is replaced by $\mathbb{I}_{s_i=s_i'}/2 + \alpha_i^*(s_i'|s_i)/2$ and $\gamma^*(s)$ by $\gamma^*(s)/2$. Thus, without loss, we can assume $\alpha_i^*(s_i|s_i) > 0$ for every s_i, which implies that α_i^* is aperiodic. Thus, $\lim_{n\to\infty}(\alpha_i^*)^n$ exists (Stroock (2014), equation (4.1.15)), and let

$$P_i = \lim_{n\to\infty}(\alpha_i^*)^n. \tag{6}$$

The α^*-reduction of the information structure (S,σ) is the information structure $(\widehat{S},\widehat{\sigma})$ for which $\widehat{S}_i$ is the set of recurrent classes of α_i^* (so $\widehat{s}_i$ is a subset of S_i for each $\widehat{s}_i \in \widehat{S}_i$, and all $s_i \in \widehat{s}_i$ has the same $P_i(\cdot|s_i)$), and

$$\widehat{\sigma}(\widehat{s},\theta) = \sum_s \sigma(s,\theta)\prod_i P_i(\widehat{s}_i|s_i),$$

where $P_i(\widehat{s}_i|s_i) = \sum_{s_i'\in\widehat{s}_i} P_i(s_i'|s_i)$. That is, $\widehat{\sigma}$ is an individual garbling of σ (cf. Brooks, Du, and Haberman (2024)).

Proposition 3. *Let α^* be part of a solution to the dual program for (S,σ), and suppose $\alpha_i^*(s_i|s_i) > 0$ for every i and s_i. Let $(\widehat{S},\widehat{\sigma})$ be the α^*-reduction. For any mechanism (A,m) and equilibrium $\widehat{\beta}$ of $(m,\widehat{\sigma})$, the following strategies are an equilibrium of (m,σ):*

$$\beta_i(a_i|s_i) = \sum_{\widehat{s}_i} P_i(\widehat{s}_i|s_i)\widehat{\beta}_i(a_i|\widehat{s}_i), \tag{7}$$

where P_i is defined in (6). *Moreover, we have $W(\beta;m,\sigma) = W(\widehat{\beta};m,\widehat{\sigma})$.*

Proof. We will prove that β is an equilibrium of (m,σ). The second part follows by construction.

Fix an agent i, and for any strategy $\beta_i' : S_i \to \Delta(A_i)$, by (5) we have

$$\left(\sum_a m(\omega|a)\beta_i'(a_i|s_i)\prod_{j\neq i}\beta_j(a_j|s_j)\right)\left(\sum_{j,\theta} u_j(\omega,\theta)\left(\sum_{s_j'}\alpha_j^*(s_j|s_j')\sigma(s_j',s_{-j},\theta) - \sigma(s,\theta)\right) + \gamma^*(s)\right) \geqslant 0,$$

for every s and ω.

Summing over s and ω, the above equation becomes

$$\sum_{\omega,\theta,a,s} u_i(\omega,\theta)\left(\sum_{s_i'}\alpha_i^*(s_i|s_i')\sigma(s_i',s_{-i},\theta)\beta_i'(a_i|s_i) - \sigma(s,\theta)\beta_i'(a_i|s_i)\right)m(\omega|a)\prod_{j\neq i}\beta_i(a_j|s_j) \geqslant 0,$$

since $\sum_{s'_j, s_j} \alpha^*_j(s_j|s'_j)\sigma(s'_j, s_{-j}, \theta)\beta_j(a_j|s_j) - \sum_{s_j} \sigma(s,\theta)\beta_j(a_j|s_j) = 0$ and $\sum_s \gamma^*(s) = 0$.

Thus, we have

$$\sum_{\omega,\theta,a,s} u_i(\omega,\theta) \left(\sum_{s'_i} (\alpha^*_i)^n(s_i|s'_i)\sigma(s'_i, s_{-i}, \theta)\beta'_i(a_i|s_i) - \sigma(s,\theta)\beta'_i(a_i|s_i) \right) m(\omega|a) \prod_{j\neq i} \beta_i(a_j|s_j) \geqslant 0$$

for every $n \geqslant 1$. If we take β'_i to be a best response of β_{-i} in (m,σ) and send $n \to \infty$, we conclude that β_i must be a best response to β_{-i} in (m,σ). □

Thus, dual reduction effectively shrinks the set of equilibrium outcomes of the information structure.

As in Section 4, we can iteratively reduce the information structure (S,σ). At each round, as long as $\alpha^*_i(s'_i|s_i) > 0$ for some $s'_i \neq s_i$, the number of recurrent classes of α^*_i is strictly less than $|S_i|$. Thus, every round of non-trivial dual reduction decreases the number of signals of some agent by at least one. This proves that there can be at most finitely many rounds of dual reduction.

When iterated dual reduction converges to an information structure $(\widehat{S},\widehat{\sigma})$, it must be that for that mechanism, the unique optimal dual solution must be $\alpha^*_i(s_i|s_i) = 1$. By strong complementary slackness in linear programming (Goldman and Tucker, 1956, Corollary 2B), there must exist a solution of the primal program for which all incentive compatibility constraints are strict. We obtain the following result:

Proposition 4. *For any information structure (S,σ), there exists an elementary information structure $(\widehat{S},\widehat{\sigma})$, which is obtained from (S,σ) via iterative dual reduction. For any mechanism (A,m) and equilibrium $\widehat{\beta}$ of $(m,\widehat{\sigma})$, there is an equilibrium β of (m,σ) with $W(\beta;m,\sigma) = W(\widehat{\beta};m,\widehat{\sigma})$.*

As in our discussion of dual reduction of mechanisms, we have thus far not mentioned any kind of participation constraint. What is the relationship between dual reduction and interim individual rationality? Consider the following example. There are two states of the world, high and low, which are both equally likely ex ante. In the high state, Ann gets a payoff of 1 from every outcome, and in the low state she gets a payoff of -1 from every outcome. Thus, no matter the information structure, Ann is always indifferent between all actions, and hence every information structure dually reduces to a trivial one in which Ann has no information. Moreover, if Ann has full information about the state, then no mechanism and equilibrium would satisfy interim individual rationality: in the low state, Ann would always strictly prefer her outside option. But under no information, Ann gets an expected payoff of 0 from every outcome, and is therefore willing to participate.

This example shows that while dual reduction shrinks the set of equilibrium outcomes, mechanism by mechanism, it may *expand* the set of equilibrium outcomes that are interim individually rational.

Now, in a mechanism that is strongly individually rational, all equilibrium outcomes in all information structures are interim individually rational. (Note that in the preceding pathological example, no strongly individually rational mechanism exists.) Hence, dual reduction of information structures shrinks the set of interim individually rational equilibrium outcomes for every strongly individually rational mechanism, and a fortiori the same is true for participation secure mechanisms.

This observation is useful in applying dual reduction to information design, as we now explain. The *potential* of an information structure (S, σ), as introduced by Brooks and Du (2024b), is defined to be:

$$P(\sigma) = \sup_{(A,m)\in\mathcal{M}_0} \sup_{\beta\in E(m,\sigma)} W(\beta; m, \sigma),$$

that is, the highest expected welfare on (S, σ) across all participation secure mechanisms and equilibria.[8] An immediate corollary of Proposition 4 is:

Corollary 2. *Let $\widehat{\mathcal{I}}(\mu)$ be the set of elementary information structure compatible with a prior $\mu \in \Delta(\Theta)$. Then we have*

$$\inf_{(S,\sigma)\in\mathcal{I}(\mu)} P(\sigma) = \inf_{(S,\sigma)\in\widehat{\mathcal{I}}(\mu)} P(\sigma).$$

Bergemann, Brooks, and Morris (2016) and Brooks and Du (2021) solve for the information structures that minimize the potential in a common value auction. Those papers solve for the min potential by relaxing participation security to interim individual rationality. We refer to the value of this relaxed program as the *weak potential.* The same approach is used in Brooks and Du (2023, 2024b) and Brooks, Du, and Feffer (2026a).

The weak potential is always weakly larger than the potential as we have defined it here, but in these applications, the minimum potential is equal to minimum weak potential.

In certain applications, there are outcomes in the mechanism that are payoff equivalent to the outside option for all agents. In the auction context, this could be withholding the good and no transfers to or from the agents. More generally, suppose there is an outcome ω^* for which $u_i(\omega^*, \theta) = 0$ for all i and θ. Consider any information structure,

[8] We might have defined the potential to be optimized over all strongly individually rational mechanisms. In principle this could lead to a different and strictly higher min potential, but it does not for the applications cited below, and Corollary 2 would still be true as stated.

mechanism, and equilibrium for which interim individual rationality is satisfied. Then the designer could enrich the mechanism by adding a secure action for each agent, by which they may unilaterally cause ω^* to be implemented with probability one. This enriched mechanism is participation secure. But clearly, no agent has an incentive to deviate to this explicit outside option. So, the original strategies are still an equilibrium of the enriched mechanism, with the same payoff and welfare.

Our point is that in environments with an explicit outside option, participation security imposes no additional constraints on the potential beyond interim individual rationality, and the potential is equal to the weak potential. Of course, having an explicit outside option is sufficient but not necessary for the two to coincide: Brooks and Du (2021) also solve for the potential minimizing information structure in the case where the designer has to allocate the good with probability one. In this case, there is no explicit outside option, since in some states every implementable outcome involves some buyer getting a higher payoff than their outside option.

As shown in Brooks and Du (2024a), the program of minimizing the weak potential (which is simply called the potential in that paper) can be simplified in an analogous manner to the program of guarantee maximization: It is without loss to restrict attention to information structures in which each agent's signals are linearly ordered, the lowest type has a binding interim individual rationality constraint, and the only relevant equilibrium constraints are those that are local and pointing towards the lowest type. Moreover, it is without loss to normalize the units for signals so that the non-zero Lagrange multipliers on all of the participation and equilibrium constraints have the same value. This effectively reduces the task of minimizing the weak potential to solving a one-dimensional family of linear programming problems, parametrized by the Lagrange multiplier. We do not know of any analogous result for minimizing the potential, as defined here over all participation secure mechanisms. Brooks and Du (2024a) give an example for which the minimum weak potential is strictly larger than the minimum potential.

We have not yet satisfactorily resolved the mystery of when the minimum weak potential is equal to the minimum potential. We have also not yet resolved another major open question in this theory: Whether and when the min potential is equal to the max guarantee. This has turned out to be the case in all of the substantive applications of the theory that have been solved so far, and Brooks and Du (2024b) show that the min potential, min weak potential, and max guarantee all coincide for a class of auction problems. It would not surprise us if a deeper understanding of the relationship between the potential, weak potential, and the guarantee, is achieved simultaneously through some novel insight about the structure of the model, or if dual reduction turns out to play an important role in that part of the story. But that we leave for a future installment in this research agenda.

References

Abreu, D. and H. Matsushima (1992): "Virtual Implementation in Iteratively Undominated Strategies: Complete Information," *Econometrica*, 993–1008.

Bergemann, D., B. Brooks, and S. Morris (2016): "Informationally Robust Optimal Auction Design," Tech. rep., Princeton University and the University of Chicago and Yale University, working paper.

Bergemann, D. and S. Morris (2005): "Robust mechanism design," *Econometrica*, 1771–1813.

——— (2013): "Robust Predictions in Games with Incomplete Information," *Econometrica*, 81, 1251–1308.

——— (2016): "Bayes Correlated Equilibrium and the Comparison of Information Structures in Games," *Theoretical Economics*, 11, 487–522.

Brooks, B. and S. Du (2021): "Optimal auction design with common values: An informationally robust approach," *Econometrica*, 89, 1313–1360.

——— (2023): "Robust Mechanisms for the Financing of Public Goods," Tech. rep., The University of Chicago and University of California-San Diego, working paper.

——— (2024a): "Dual Reductions and the First-Order Approach for Informationally Robust Mechanism Design," *University of Chicago, Becker Friedman Institute for Economics Working Paper*.

——— (2024b): "On the structure of informationally robust optimal mechanisms," *Econometrica*, 92, 1391–1438.

——— (2025): "Simplicity and Portability in Mechanism Design: A Case for (and Against) the Worst Case," .

Brooks, B., S. Du, and J. Feffer (2026a): "Compound Proportional Auctions," Working paper, The University of Chicago and University of California-San Diego and Stanford University.

Brooks, B., S. Du, and A. Haberman (2024): "Robust Predictions with Bounded Information," Tech. rep., The University of Chicago and University of California-San Diego and Stanford University, working paper.

Brooks, B., S. Du, and L. Zhang (2026b): "Market Order Mechanisms," Working paper, The University of Chicago and University of California-San Diego.

Che, E. (2022): "Robustly optimal auction design under mean constraints," in *Proceedings of the 23rd ACM Conference on Economics and Computation*, 153–181.

Chen, Y.-C. and J. Li (2018): "Revisiting the Foundations of Dominant-Strategy Mechanisms," *Journal of Economic Theory*, 178, 294–317.

Chung, K.-S. and J. C. Ely (2007): "Foundations of Dominant-Strategy Mechanisms," *The Review of Economic Studies*, 74, 447–476.

Dasgupta, P., P. Hammond, and E. Maskin (1979): "The implementation of social choice rules: Some general results on incentive compatibility," *The Review of Economic Studies*, 46, 185–216.

Du, S. (2018): "Robust mechanisms under common valuation," *Econometrica*, 86, 1569–1588.

Goldman, A. J. and A. W. Tucker (1956): "Theory of Linear Programming," in *Linear Inequalities and Related Systems*, ed. by H. W. Kuhn and A. W. Tucker, Princeton, NJ: Princeton University Press, 53–98.

Haberman, A. and R. Jagadeesan (2025): "Auctions with Withdrawal Rights: A Foundation for Uniform Price," in *Proceedings of the 26th ACM Conference on Economics and Computation*, 35–35.

He, W. and J. Li (2022): "Correlation-robust auction design," *Journal of Economic Theory*, 200, 105403.

He, W., J. Li, and W. Zhong (2024): "Rank-guaranteed auctions," *arXiv preprint arXiv:2408.12001*.

Maskin, E. (1999): "Nash Equilibrium and Welfare Optimality," *The Review of Economic Studies*, 66, 23–38.

Milgrom, P. R. and R. J. Weber (1982): "A Theory of Auctions and Competitive Bidding," *Econometrica*, 1089–1122.

Myerson, R. (2024): "Dual Reduction and Elementary Games with Senders and Receivers," Tech. rep., The University of Chicago, working paper.

Myerson, R. B. (1997): "Dual reduction and elementary games," *Games and Economic Behavior*, 21, 183–202.

Stroock, D. W. (2014): *An Introduction to Markov Processes*, Springer, 2 ed.

Vickrey, W. (1961): "Counterspeculation, Auctions, and Competitive Sealed Tenders," *The Journal of Finance*, 16, 8–37.

Yamashita, T. (2018): "Revenue guarantees in auctions with a (correlated) common prior and additional information," Tse working paper.

Yamashita, T. and S. Zhu (2018): "On the foundations of ex post incentive compatible mechanisms," *American Economic Journal: Microeconomics*.