Dual Reductions and the First-Order Approach for Informationally Robust Mechanism Design*

Benjamin Brooks Songzi Du

October 19, 2024

Abstract

The *guarantee* of a mechanism is the lowest objective value for the designer, across all information structures and equilibria. Brooks and Du (2024) proposed a "first-order" approach to characterizing guarantee-maximizing mechanisms by maximizing a particular lower bound on the guarantee: the expected lowest *strategic virtual objective*. In this paper, we show that for any mechanism M, there is an associated "dual reduction" mechanism M' for which the expected lowest strategic virtual objective of M' (and hence the guarantee of M') is greater than the guarantee of M. This provides a rigorous foundation for the use of the strategic virtual objective in designing informationally robust mechanisms. A parallel result, based on dual reductions of information structures, justifies Brooks and Du's (2024) first-order approach to characterizing information structures with the lowest *potential*, in terms of those that minimize the expected highest *informational virtual objective*.

KEYWORDS: Mechanism design, information design, dual reduction, max-min, Bayes correlated equilibrium, robustness.

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. This research has been supported by NSF #2215259 and #2215475.

1 Introduction

The literature on informationally-robust mechanism design has shifted the focus from a mechanism's performance in a single information environment to, instead, bounding the mechanism's performance across a large family of environments. This reflects a concern that in practical settings, a mechanism designer may not know the detailed structure of agents' information. Moreover, the standard modeling devices for distributed information often seem more plausible as loose metaphors, that could discipline our theoretical search for desirable institutions, rather than as a literal description of agents' knowledge and thought processes that should be "hard-coded" into the way agents interact with a mechanism.

In a series of recent papers, we and our coauthors have been developing a methodology for designing informationally-robust mechanisms. Holding fixed the distribution of payoff-relevant fundamentals in the economy, we define the *guarantee* of a mechanism to be the lowest expected objective for the designer, across all informational environments and equilibria. Guarantee maximizing mechanisms have been characterized in applications involving optimal auctions, bilateral trade, and public goods (Brooks and Du, 2021b,a, 2023, 2021a, 2024; Brooks, Du, and Zhang, 2024b; Brooks, Du, and Haberman, 2024a). These novel mechanisms have both strong welfare guarantees and relatively simple structure, suggesting both that they may be practically useful and that guarantee-maximization is a useful heuristic for discovering new mechanisms with desirable incentive properties.

These applications employ a general methodology that is fully described in Brooks and Du (2024): Instead of maximizing the guarantee directly, we maximize a lower bound on the guarantee, which is the expected lowest *strategic virtual objective*. This is the designer's innate objective, plus extra terms that correspond to a simple class of "local inward" equilibrium constraints, pointing away from the action that corresponds to opting out of the mechanism. The focus on this particular class of equilibrium constraints represents a kind of "first order" approach in informationally robust mechanism design, wherein we work with bounds on performance derived from first-order conditions for optimality of agents' behavior. In a series of applications, the lower bound has been found to be tight, in the sense that the maximum guarantee is equal to the maximum expected lowest strategic virtual objective.[1] In this paper, we provide a general proof that the lower bound is *always* tight. We thus provide a solid foundation for the use of the strategic virtual objective in informationally robust mechanism design.

The approach is as follows: We start from a given mechanism, and its associated guarantee. This guarantee is the optimal value of a linear program in which we minimize the designer's objective over all *Bayes correlated equilibria* (BCE) (Bergemann and Morris, 2016). The dual of this linear program involves optimization over Lagrange multipliers on a collection of "obedience" constraints: each agent prefers their equilibrium action over any alternative. Up to rescaling, these multipliers can be interpreted as mixed strategies that

[1]It is important to note that in classical mechanism design, strong assumptions are needed on the informationally environment for local equilibrium constraints to pin down the optimal mechanism. In contrast, the guarantee holds across *all* information structures and equilibria, including those with much more complicated incentive structure. Nonetheless, in the aforementioned applications, it turns out that the critical "worst-case" information structures are those for which a simple pattern of binding equilibrium constraints emerges at the optimum.

represent the most "tempting" deviation from an equilibrium pure action. In particular, the optimal value of the linear program would not change if instead of imposing all obedience constraints, we only impose the constraint that an agent not want to deviate to this most tempting mixture. We use these most tempting deviations to construct a new "reduced" mechanism, in which there are countably many pure actions, each of which corresponds to a particular mixed strategy in the original mechanism. Specifically, the first reduced action corresponds to opting out (for sure) in the original mechanism. This is followed by a second reduced action, which corresponds to the most tempting deviation from opting out. The third reduced action corresponds to the mixture that would result by drawing from the second reduced action, and then following the most tempting deviation from what was drawn. And so on, thus generating a countable sequence of reduced actions, where each element in the sequence corresponds to drawing from its predecessor and then taking the most tempting deviation. We then show that the expected lowest strategic virtual objective for the reduced mechanism is precisely the optimal Lagrangian that pins down the guarantee in the original mechanism.

This construction represents a kind of "dual reduction" of the original mechanism that is analogous to, but distinct from, that of Myerson (1997). We discuss this connection in more detail in Section 5.

It is important to note that our construction starts with an original mechanism that has finitely many actions, whereas the reduced mechanism has countably infinitely many actions. A straightforward consequence of our proof, however, is that by truncating this construction after a sufficiently large number of steps, we obtain a finite reduced mechanism whose associated lower bound is arbitrarily close to the guarantee of the original mechanism.

How can we be certain that a given mechanism maximizes the guarantee? Brooks and Du (2024) adopt the following duality-based approach. The *potential* of an information structure is the highest expected objective of the designer, across all mechanisms and equilibria. It is straightforward to establish that the minimum potential must be at least the maximum guarantee. A saddle point for the guarantee/potential problem would be a pair of a mechanism and an information structure, such that the respective guarantee and potential are equal. By standard arguments, if a saddle point exists, then its mechanism maximizes the guarantee and its information structure minimizes the potential.

Analogous to their approach to maximizing the guarantee, Brooks and Du (2024) also provided a "first order" approach to minimizing the potential. In particular, we proved that an information structure's potential is less than its expected highest *informational virtual objective*. The latter is the designer's innate objective, plus additional terms that correspond to a particular pattern of equilibrium constraints: the signals are linearly ordered, and binding equilibrium constraints are "local outward" pointing towards a single signal with a binding participation constraint. In this paper, we provide a proof that the potential upper bound is also tight. The construction shares a high level similarity with that described above: Given an information structure, we construct a "dual reduction" information structure, in which the agents' signals are essentially mixtures over signals in the original information structure. These mixtures are defined from the optimal Lagrange multipliers for the potential-minimization program, which correspond to either most tempting distributions of signals to mimic, or probabilistic deviations to opting out. We then show that the expected highest informational virtual objective for the reduced information structure is less than the potential of the original information structure. This dual reduction involves countably infinitely many signals, but again, we also establish an approximate version of our result with a finite dual reduction.

The rest of this paper is as follows. Section 2 introduces notation and terminology. Section 3 presents our result on dual reductions of mechanisms. Section 4 presents the dual analysis for dual reductions of information structures. Section 5 is a discussion and conclusion.

2 Model

There is a mechanism designer and a finite group of agents indexed by $i \in \{1, \ldots, N\}$. The designer controls an outcome $\omega \in \Omega$, where Ω is finite. The designer and the agents have expected utility preferences over outcomes. In particular, the preferences of agent $i = 1, \ldots, N$ over outcomes and states are represented by the utility index $u_i(\omega, \theta)$, which depends on a payoff-relevant state of the world $\theta \in \Theta$, where Θ is also finite. The designer's preferences are similarly represented by the utility index $w(\omega, \theta)$. The designer has a prior belief about θ, denoted $\mu \in \Delta(\Theta)$, which is held fixed throughout our analysis.[2]

Each agent could choose not to participate in the designer's mechanism and receive a certain state-dependent payoff. We normalize this outside option to zero and interpret agent i's utility as their payoff net of the outside option.

The agents' private information about θ is described by an *information structure*, which consists of: a product set of signal profiles $S = \prod_i S_i$,[3] where S_i is agent i's set of signals, and a joint distribution $\sigma \in \Delta(S \times \Theta)$ for which the marginal on Θ is μ. We assume that the S_i are finite or countably infinite.[4] The information structure is finite if each S_i is finite. An information structure is denoted $I = (S, \sigma)$, $\mathcal{I}$ is the set of information structures, and $\overline{\mathcal{I}}$ is the subset of information structures that are finite.[5]

The designer commits to a *mechanism*, which consists of: a product set of action profiles $A = \prod_i A_i$, where A_i is agent i's set of actions, and an outcome function $m : A \to \Delta(\Omega)$ that maps action profiles to lotteries over outcomes. We assume that A_i is finite or countably infinite. The mechanism is finite if each A_i is finite. An action $a_i \in A_i$ is *participation secure* if $\sum_\omega u_i(\omega, \theta) m(\omega | a_i, a_{-i}) \geqslant 0$ for all a_{-i} and θ. A mechanism is *participation secure* if every agent has an action that is participation secure. We will restrict the mechanism designer to use only mechanisms that are participation secure. This ensures that, regardless of the information structure and other agents' strategies, no agent will have a strict incentive to

[2]Portions of this section are replicated almost verbatim from Section 2 of Brooks and Du (2024).

[3]Throughout our exposition, a sum or a product with respect to a variable without qualification means that the operation should be applied for all values of the variable. In this case, the product is over all i, that is, $i = 1, \ldots, N$.

[4]Brooks and Du (2024) restrict attention to finite mechanisms and information structures (cf. the discussion in Section 3.3, *ibid*). We allow countably infinity here so that we can state exact results for the dual reductions, which are countably infinite, and we report analogous approximate results when we restrict to finite objects.

[5]The set of finite or countably infinite information structures is defined by identifying finite sets of signals with finite subsets of the natural numbers. Likewise for the set of mechanisms.

exit the mechanism, since they can always play a participation secure action and receive a weakly higher payoff than their outside option. A mechanism is denoted by $M = (A, m)$, the set of all mechanisms is $\mathcal{M}$, the set of participation secure mechanisms is $\mathcal{M}^*$, and the set of finite participation secure mechanisms is $\overline{\mathcal{M}}^*$. We assume that a participation secure mechanism exists.

A mechanism and an information structure (M, I) together define a *Bayesian game*, in which a *(behavioral) strategy* for agent i is a mapping $b_i : S_i \to \Delta(A_i)$. A strategy profile $b = (b_1, \ldots, b_N)$ is identified with the function from S to $\Delta(A)$ defined by $b(a|s) = \prod_i b_i(a_i|s_i)$. Expected utility for agent i is

$$U_i(M, I, b) = \sum_{\theta,s,a,\omega} u_i(\omega, \theta)m(\omega|a)b(a|s)\sigma(s, \theta),$$

and the designer's welfare is

$$W(M, I, b) = \sum_{\theta,s,a,\omega} w(\omega, \theta)m(\omega|a)b(a|s)\sigma(s, \theta).$$

A strategy profile b is a *(Bayes Nash) equilibrium* of (M, I) if $U_i(M, I, b) \geqslant U_i(M, I, b'_i, b_{-i})$ for all $i = 1, \ldots, N$ and b'_i. The set of equilibria is $\mathcal{E}(M, I)$, which we note is non-empty whenever M and I are both finite.

The *guarantee* of a mechanism M is

$$G(M) = \inf_{I \in \overline{\mathcal{I}}} \inf_{b \in \mathcal{E}(M,I)} W(M, I, b),$$

that is, the infimum welfare of the designer across all information structures and equilibria. The *potential* of an information structure I is

$$P(I) = \sup_{M \in \overline{\mathcal{M}}^*} \sup_{b \in \mathcal{E}(M,I)} W(M, I, b),$$

that is, the supremum welfare of the designer across all participation-secure mechanisms and equilibria. It is immediate that for any $M \in \overline{\mathcal{M}}^*$ and $I \in \overline{\mathcal{I}}$, $G(M) \leqslant P(I)$.

3 Dual reductions of Mechanisms

This section presents our first main result: For any mechanism M, we can construct a "dual reduction" of it, denoted M', for which the expected lowest strategic virtual objective of M' is greater than $G(M)$. To develop this result, we first describe the strategic virtual objective, and then we will construct the dual reduction.

The strategic virtual objective is defined by first fixing a linear order on actions. Without loss, we label the actions with integers. Let $X_i^* = \{0, 1, 2, \ldots\}$ to be the non-negative integers, and $X^* = \prod_i X_i^*$. We will treat X^* as an action space and consider *ordered mechanisms* of the form (X^*, m). $\mathcal{M}^*$ is the set of ordered mechanisms for which $0 \in X_i^*$ is participation secure for each agent i.

In a slight abuse of notation, we say that an ordered mechanism (X^*, m) is finite if there is some $k \in X_i^*$ such that for i, ω, and x with $x_i > k$, $m(\omega|x) = m(\omega|k, x_{-i})$. Thus, all actions above some finite threshold have the same meaning in the mechanism.

Fix an ordered mechanism $(X^*, m) \in \mathcal{M}^*$ and a constant $C \in \mathbb{R}_+$. The associated *strategic virtual objective* at the action profile x and the state θ is

$$\lambda(x, \theta, C) \equiv \sum_\omega \left[w(\omega, \theta) m(\omega|x) + C \sum_i u_i(\omega, \theta)(m_i(\omega|x_i + 1, x_{-i}) - m_i(\omega|x)) \right].$$

As we mentioned in the introduction, λ is the sum of the designer's innate objective, plus terms that correspond to the agents' gains from deviating to actions which are further away from the participation secure action. The constant C is effectively the multiplier on the agents' equilibrium constraints. We have fixed this multiplier to be the same for all agents and equilibrium actions, but as discussed in Brooks and Du (2024, Section 3.3), this is essentially without loss.

Proposition 3 of Brooks and Du (2024) shows that for any ordered mechanism (X^*, m),

$$G(X^*, m) \geqslant \underline{G}(X^*, m, C) \equiv \sum_\theta \mu(\theta) \inf_x \lambda(x, \theta, C).$$

Now, starting from a finite mechanism $M = (A, m)$, we define a dual reduction $(X^*, m^*) \in \mathcal{M}^*$ as follows. Let σ^* be a BCE that attains the minimum in $G(M)$, and let $\alpha_i^*(a_i'|a_i)$ be the associated optimal multipliers on obedience constraints. Thus,[6]

$$\begin{aligned} G(M) = &\min_{\substack{\sigma \in \Delta(A \times \Theta) \\ \text{s.t. } \mathrm{marg}_\Theta\, \sigma = \mu}} \sum_{a, \theta, \omega} \sigma(a, \theta) \left[w(\omega, \theta) m(\omega|a) + \sum_{i, a_i'} \alpha_i^*(a_i'|a_i) u_i(\omega, \theta)(m(\omega|a_i', a_{-i}) - m(\omega|a_i, a_{-i})) \right] \\ &= \sum_\theta \mu(\theta) \min_a \sum_\omega \left[w(\omega, \theta) m(\omega|a) + \sum_{i, a_i'} \alpha_i^*(a_i'|a_i) u_i(\omega, \theta)(m(\omega|a_i', a_{-i}) - m(\omega|a_i, a_{-i})) \right]. \end{aligned}$$

Let

$$C = 1 + \max_{a_i} \sum_{a_i' \neq a_i} \alpha_i^*(a_i'|a_i),$$

and without loss, set

$$\alpha_i^*(a_i|a_i) = C - \sum_{a_i' \neq a_i} \alpha_i^*(a_i'|a_i) > 0.$$

Note that we have intentionally used the same notation for the normalizing constant as for the Lagrange multiplier on local constraints. In the proof of Theorem 1 below, they will turn out to be the same quantity.[7]

[6] The Lagrangian in $G(M)$ would coincide with the strategic virtual objective of M if $A_i = X_i^*$ and $\alpha_i^*(s_i'|s_i) = C$ if $a_i' = a_i + 1$ and is zero otherwise.

[7] Indeed, this construction provides a complementary perspective as to why it is without loss to normalize the "size" of a deviation so that the multiplier is constant: we are free to choose the probability $\alpha_i^*(a_i|a_i)$, which is proportional to the likelihood of not deviating at all, and parametrizes the likelihood of a deviation from the given action

For each i, let a_i^0 be (any) participation secure action for player i.[8] Then, for each i, define a function $g_i : X_i^* \to \Delta(A_i)$ as follows: $g_i(a_i^0|0) = 1$, and for $k > 0$,

$$g_i(a_i|k) = \sum_{a_i'} g_i(a_i'|k-1) \left[\frac{\alpha_i^*(a_i|a_i')}{\sum_{a_i''} \alpha_i^*(a_i''|a_i')} \right].$$

The interpretation is as follows: each action $x_i \in X_i^*$ corresponds to a mixture in the original mechanism. The action $x_i = 0$ is identified with $a_i^0 \in A_i$, a participation secure action; $x_i = 1$ is the most tempting deviation from a_i^0. Action x_i is the mixture obtained by drawing an action from $x_i - 1$ and then taking the most tempting deviation from the corresponding original action. In other words, x_i is the mixture obtained by iteratively taking the most tempting deviation x_i times, starting from a_i^0. Thus, actions in the dual reduction are interpreted as a number of deviations from a_i^0.

Finally, we define for all ω and x,

$$m^*(\omega|x) = \sum_{a \in A} g(a|x) m(\omega|a),$$

where

$$g(a|x) = \prod_i g_i(a_i|x_i).$$

Theorem 1. *For any finite participation secure mechanism $M = (A, m)$ and corresponding dual reduction (X^*, m^*), we have that $\underline{G}(X^*, m^*, C) \geqslant G(M)$.*

As a result, the supremum guarantee across finite participation secure mechanisms is equal to the supremum expected lowest strategic virtual objective across all participation secure ordered mechanisms.[9]

Proof. We have:

$$\begin{aligned}
&\underline{G}(X^*, m^*) \\
&= \sum_\theta \mu(\theta) \inf_x \sum_\omega \left[w(\omega,\theta) m^*(\omega|x) + C \sum_i (m^*(\omega|x_i+1, x_{-i}) - m^*(\omega|x_i, x_{-i})) u_i(\omega,\theta) \right] \\
&= \sum_\theta \mu(\theta) \inf_x \sum_{\omega,a} \left[w(\omega,\theta) m(\omega|a) g(a|x) + C \sum_i m(\omega|a) \left(g_i(a_i|x_i+1) - g_i(a_i|x_i) \right) g_{-i}(a_{-i}|x_{-i}) u_i(\omega,\theta) \right] \\
&= \sum_\theta \mu(\theta) \inf_x \sum_{\omega,a} \Bigg[w(\omega,\theta) m(\omega|a) g(a|x) \\
&\qquad\qquad + C \sum_i m(\omega|a) \left(\sum_{a_i'} \frac{\alpha_i^*(a_i|a_i')}{C} g_i(a_i'|x_i) - g_i(a_i|x_i) \right) g_{-i}(a_{-i}|x_{-i}) u_i(\omega,\theta) \Bigg]
\end{aligned}$$

[8]There is a different dual reduction for every choice of optimal multipliers, every selection of the participation secure actions, and every choice of the constant C.

[9]The latter supremum is taken across all participation secure ordered mechanisms and multipliers C.

$$= \sum_\theta \mu(\theta) \inf_x \sum_{\omega,a} \left[w(\omega,\theta)m(\omega|a)g(a|x) + \sum_{i,a_i'} \alpha_i^*(a_i'|a_i) \left(m(\omega|a_i', a_{-i}) - m(\omega|a_i, a_{-i})\right) g(a|x) u_i(\omega,\theta) \right].$$

The above equation is clearly weakly larger than

$$\sum_\theta \mu(\theta) \min_a \sum_\omega \left[w(\omega,\theta)m(\omega|a) + \sum_{i,a_i'} \alpha_i^*(a_i'|a_i) \left(m(\omega|a_i', a_{-i}) - m(\omega|a_i, a_{-i})\right) u_i(\omega,\theta) \right] = G(A,m).$$

□

The dual reduction that we constructed before Theorem 1 involves countably infinitely many actions. In contrast, the mechanism that we started with has only finitely many actions. An advantage of working with finite mechanisms and information structures is that equilibria always exist, so that we can be assured that favorable guarantees are not relying on some controversial use of equilibrium existence in infinite games. In fact, we can provide an approximate version of Theorem 1 with finite mechanisms.

Given a dual reduction (X^*, m^*), we define its *k-truncation* to be the mechanism $(X^*, m^* \circ f^k)$, where $f^k : X^* \to X^*$ is defined by

$$f_i^k(x) = \min(x_i, k).$$

In other words, all actions above k are relabeled as k. Clearly, $(X^*, m^* \circ f^k)$ is a finite mechanism.

Proposition 1. *For any finite participation secure mechanism $M = (A, m)$ and corresponding dual reduction (X^*, m^*) and for any $\epsilon > 0$, there exists a k so that $\underline{G}(X^*, m^* \circ f^k) \geqslant G(M) - \epsilon$.*

As a result, the supremum guarantee across finite participation secure mechanisms is equal to the supremum expected lowest strategic virtual objective across all finite participation secure ordered mechanisms.

Proof. The analogue of g for $(X^*, m^* \circ f^k)$ is $\widehat{g}_i(a_i|x_i) = g(a_i|f^k(x_i))$. Let the associated mechanism be denoted by $\widehat{m} = m^* \circ f^k$:

$$\widehat{m}(\omega|x) = \sum_{a \in A} \widehat{g}(a|x) m(\omega|a).$$

We compare the strategic virtual objectives of (X^*, m^*) and $(X^*, \widehat{m})$:

$$\begin{aligned} \lambda^*(x,\theta,C) &= \sum_\omega \left[w(\omega,\theta)m^*(\omega|x) + C \sum_i (m^*(\omega|x_i+1, x_{-i}) - m^*(\omega|x_i, x_{-i})) u_i(\omega,\theta) \right] \\ &= \sum_{\omega,a} \left[w(\omega,\theta)m(\omega|a)g(a|x) + C \sum_i m(\omega|a) \left(g_i(a_i|x_i+1) - g_i(a_i|x_i)\right) g_{-i}(a_{-i}|x_{-i}) u_i(\omega,\theta) \right], \end{aligned}$$

$$\widehat{\lambda}(x,\theta,C) = \sum_{\omega}\left[w(\omega,\theta)\widehat{m}(\omega|x) + C\sum_i(\widehat{m}(\omega|x_i+1,x_{-i}) - \widehat{m}(\omega|x_i,x_{-i}))u_i(\omega,\theta)\right]$$

$$= \sum_{\omega,a}\left[w(\omega,\theta)m(\omega|a)\widehat{g}(a|x) + C\sum_i m(\omega|a)\left(\widehat{g}_i(a_i|x_i+1) - \widehat{g}_i(a_i|x_i)\right)\widehat{g}_{-i}(a_{-i}|x_{-i})u_i(\omega,\theta)\right],$$

$$|\lambda^*(x,\theta,C) - \widehat{\lambda}(x,\theta,C)|$$
$$\leqslant M\sum_{\omega,a}|g(a|x) - \widehat{g}(a|x)| + CM\sum_{\omega,a,i}(|g_i(a_i|x_i+1) - g_i(a_i|x_i)|\mathbb{I}_{x_i\geqslant k} + |g_{-i}(a_{-i}|x_{-i}) - \widehat{g}_{-i}(a_{-i}|x_{-i})|),$$

where M is a constant such that $|u_i(\omega,\theta)| \leqslant M$ and $|w(\omega,\theta)| \leqslant M$ for all ω and θ.

Because $\alpha_i^*(a_i|a_i) > 0$ for every a_i, every a_i is aperiodic in the Markov chain $\alpha_i^*(a_i'|a_i)/C$. Therefore, by a standard result on Markov chain (e.g., Stroock, 2014, equation (4.1.15) on page 85), $\lim_{k\to\infty} g_i(k)$ exists in $\Delta(A_i)$, which is the invariant measure of the chain when it starts from a_i^0. We denote this invariant measure by $g_i(\infty)$.

Since

$$|g_i(a_i|x_i) - \widehat{g}_i(a_i|x_i)| \leqslant |g_i(a_i|x_i) - \widehat{g}_i(a_i|k)|\mathbb{I}_{x_i\geqslant k}$$

and

$$\lim_{k\to\infty} g_i(a_i|k) = g_i(a_i|\infty)$$

for every a_i and x_i, we see that $\sup_{x,\theta}|\lambda^*(x,\theta,C) - \widehat{\lambda}(x,\theta,C)| \to 0$ as $k\to\infty$. Proposition 1 then follows from Theorem 1. □

In Brooks, Du, and Zhang (2024b), we computed guarantees for binary action trading mechanisms, where agents simply indicate whether or not they want to trade. We remarked that the guarantee would be the same even if we relaxed the solution concept to *coarse Bayes correlated equilibrium*, which is analogous to BCE, except that we only impose obedience constraints of the form: for every i and a_i, agent i should weakly prefer their equilibrium strategy to the strategy of always playing a_i (regardless of agent i's private information).

For any mechanism M, we could define its *coarse guarantee* $G^C(M)$ to be minimum designer welfare across all coarse Bayes correlated equilibria. This is a linear program, and in the dual program, relaxing the solution concept to coarse BCE manifests itself as a functional form restriction that for all i and a_i, the multiplier $\alpha_i(a_i|a_i')$ does not depend on the equilibrium action a_i'. In other words, the "most tempting deviation" is independent of the equilibrium action.

By following the same procedure as described at the beginning of this section, we can use the optimal multipliers to define a *coarse dual reduction* of M, which is derived from optimal multipliers. Because of the property described in the previous paragraph, the Markov chain on most tempting deviations converges after one step. Hence, a coarse dual reduction only has two actions: the participation secure action and the most tempting deviation. We therefore have the following corollary of Theorem 1.

Corollary 1. *For any participation secure mechanism M and corresponding coarse dual reduction (X^*, m^*), we have that $\underline{G}(X^*, m^*) \geqslant G^C(M)$. Moreover, the coarse dual reduction (X^*, m^*) has only two actions, in the sense that for all i and x_{-i} and $x_i > 0$, $m^*(x_i, x_{-i}) = m^*(1, x_{-i})$ (all positive actions are equivalent to the action 1).*

As a result, supremum coarse guarantee across all finite participation secure mechanisms is equal to the supremum expected lowest strategic virtual objective across binary action participation secure mechanisms.

4 Dual reductions of Information Structures

We now present a corresponding result for dual reductions of information structures.

Let $\overline{X}^* = \{0, 1, 2, \ldots\} \cup \{\infty\}$. In other words, we take X^* and add a point at infinity. We now consider *ordered information structures* of the form $(\overline{X}^*, \sigma)$. Such an information structure is *finite* if there exists a k such that for all x and θ, if $k < x_i < \infty$ for some i, then $\sigma(x, \theta) = 0$.

For any ordered information structure and constant C, we define the *informational virtual objective* at a signal x and for an outcome ω to be[10]

$$\psi(x, \omega, C) = \sum_\theta \left[w(\omega, \theta)\sigma(x, \theta) - C \sum_i u_i(\omega, \theta) \left(\sigma(x_i + 1, x_{-i}, \theta) - \sigma(x_i, x_{-i}, \theta)\right) \right].$$

Note the implicit convention that $\infty + 1 = \infty$. In effect, we drop all participation constraints except participation for one type, and local "outward" constraints that represent deviation towards the participation constraint. The one exception is the infinite type, for which no constraints bind. Proposition 3 of Brooks and Du (2024) shows that

$$P(\overline{X}^*, \sigma) \geqslant \underline{P}(\overline{X}^*, \sigma, C) \equiv \sum_x \max_\omega \psi(x, \omega, C).$$

We now describe how to construct dual reductions of a given information structure $I = (S, \sigma)$. The potential $P(I)$ is the solution to a linear program, for which there exist multipliers $\alpha_i^*(s_i'|s_i)$ and $\beta_i^*(s_i)$ such that[11]

$$P(I) = \max_{m: S \to \Delta(\Omega)} \sum_{s, \theta, \omega} \sigma(s, \theta) \Bigg[w(\omega, \theta) m(\omega|s) + \sum_i \beta_i^*(s_i) u_i(\omega, \theta) m(\omega|s) + \sum_{i, s_i'} \alpha_i^*(s_i'|s_i) u_i(\omega, \theta)(m(\omega|s) - m(\omega|s_i', s_{-i})) \Bigg]$$

[10] Brooks and Du (2024) show that for the optimal auctions problem, the expected highest informational virtual objective is minimized when signals are independent (unconditional on the state), in which case the informational virtual objective reduces to an interdependent-values "virtual value" that is familiar from auction theory (Myerson, 1981; Bulow and Klemperer, 1996). Thus, by pursuing an exercise inspired by Myerson (1997), we provide a foundation for the analysis of independent types and virtual values, following Myerson (1981).

[11] The Lagrangian in $P(I)$ would coincide with the informational virtual objective of I if $A_i = \overline{X}_i^*$ and $\alpha_i^*(s_i'|s_i) = C$ if $a_i' = a_i - 1$ and is zero otherwise, and $\beta_i^*(s_i) = C$ if $s_i = 0$ and is zero otherwise.

$$= \sum_s \max_\omega \sum_\theta \Bigg[w(\omega,\theta)\sigma(s,\theta) + \sum_i \beta_i^*(s_i) u_i(\omega,\theta)\sigma(s,\theta)$$

$$+ \sum_{i,s_i'} \left[\alpha_i^*(s_i'|s_i)\sigma(s,\theta) - \alpha_i^*(s_i|s_i')\sigma(s_i', s_{-i},\theta)\right] u_i(\omega,\theta) \Bigg].$$

To define a dual reduction of I, we associate each s_i with a distribution over $x_i \in \overline{X}_i^*$. This distribution is defined from a particular Markov chain, for which the states are elements of S_i, plus an additional absorbing state $\varnothing$. We let

$$C = 1 + \max_{s_i} \left[\beta^*(s_i) + \sum_{s_i' \neq s_i} \alpha_i^*(s_i'|s_i) \right].$$

We then set $\alpha_i^*(s_i|s_i)$ so that for all s_i,

$$C = \beta^*(s_i) + \sum_{s_i'} \alpha_i^*(s_i'|s_i).$$

Starting from s_i, with probability $\beta^*(s_i)/C$, the chain transitions to $\varnothing$. Otherwise, with probability $\alpha_i^*(s_i'|s_i)/C$, the s_i transitions to s_i'. Suppose that we draw an initial (s,θ) according to σ, and then let the Markov chain run. For each i, there is a certain number of periods x_i that the chain will run before reaching $\varnothing$. For each s_i, let $\rho_i(x_i|s_i)$ be the probability that starting at s_i, it takes x_i more periods to reach $\varnothing$. Note that $\rho(\infty|s_i) > 0$ means that s_i may transition to a recurrent class of the Markov chain without any s_i' for which $\beta_i(s_i') > 0$.

Note that $\rho_i(0|s_i) = \beta_i^*(s_i)/C$, and for $x_i > 0$, it is defined recursively as

$$\rho_i(x_i|s_i) = \sum_{s_i'} \frac{\alpha_i^*(s_i'|s_i)}{C} \rho_i(x_i - 1|s_i').$$

We further remark that it may be that the chain never transitions to zero, if s_i does not communicate with a signal s_i' for which $\beta_i(s_i') > 0$. In that case, we simply have $\rho_i(\infty|s_i) = 1$. We then define

$$\sigma^*(x,\theta) = \sum_s \sigma(s,\theta)\rho(x|s),$$

where

$$\rho(x|s) = \prod_i \rho_i(x_i|s_i).$$

In other words, $(\overline{X}^*, \sigma^*)$ is the information structure we would obtain if agents do not get to observe their original signals. Instead, we have the signals transition independently until they reach $\varnothing$, and agents observe how many periods it takes their signal to transition to $\varnothing$ (or they observe ∞ if the signal never transitions to $\varnothing$).

Theorem 2. *Given any finite information structure $I = (S, \sigma)$ and corresponding dual reduction (X^*, σ^*) and multiplier C, we have that $\underline{P}(X^*, \sigma^*, C) \leqslant P(I)$.*

As a result, the infimum potential across all finite information structures is equal to the infimum expected highest informational virtual objective across all ordered information structures.[12]

Proof. For the reduced information structure, we now have the following upper bound on designer welfare:

$$
\begin{aligned}
&\overline{P}(X^*, \sigma^*) \\
&= \sum_x \max_\omega \sum_\theta \left[w(\omega,\theta)\sigma^*(x,\theta) - C \sum_i u_i(\omega,\theta) \left(\sigma^*(x_i+1, x_{-i}, \theta) - \sigma^*(x,\theta)\right) \right] \\
&= \sum_x \max_\omega \sum_\theta \left[w(\omega,\theta)\sigma^*(x,\theta) - C \sum_s \sigma(s,\theta) \sum_i u_i(\omega,\theta) \left(\rho_i(x_i+1|s_i) - \rho_i(x_i|s_i)\right) \rho_{-i}(x_{-i}|s_{-i}) \right] \\
&= \sum_x \max_\omega \sum_{s,\theta} \Bigg[w(\omega,\theta)\sigma(s,\theta)\rho(x|s) \\
&\qquad - C\sigma(s,\theta) \sum_i u_i(\omega,\theta) \left(\sum_{s_i'} \frac{\alpha_i^*(s_i'|s_i)\rho(x_i|s_i')}{C} - \rho_i(x_i|s_i) \right) \rho_{-i}(x_{-i}|s_{-i}) \Bigg] \\
&= \sum_x \max_\omega \sum_{s,\theta} \Bigg[w(\omega,\theta)\sigma(s,\theta)\rho(x|s) \\
&\qquad - \sigma(s,\theta) \sum_i u_i(\omega,\theta) \left(\sum_{s_i'} \alpha_i^*(s_i'|s_i)\rho(x_i|s_i') - \left(\sum_{s_i'} \alpha_i^*(s_i'|s_i) + \beta_i^*(s_i) \right) \rho_i(x_i|s_i) \right) \rho_{-i}(x_{-i}|s_i) \Bigg] \\
&= \sum_x \max_\omega \sum_{s,\theta} \Bigg[w(\omega,\theta)\sigma(s,\theta)\rho(x|s) + \sum_i u_i(\omega,\theta)\beta_i^*(s_i)\sigma(s,\theta)\rho(x|s) \\
&\qquad - \sum_i u_i(\omega,\theta) \left(\sum_{s_i'} \alpha_i^*(s_i|s_i')\sigma(s_i', s_{-i}, \theta) - \sum_{s_i'} \alpha_i^*(s_i'|s_i)\sigma(s,\theta) \right) \rho(x|s) \Bigg].
\end{aligned}
$$

The above equation is clearly weakly less than

$$
\begin{aligned}
&\sum_{x,s} \max_\omega \sum_\theta \Bigg[w(\omega,\theta)\sigma(s,\theta)\rho(x|s) + \sum_i u_i(\omega,\theta)\beta_i^*(s_i)\sigma(s,\theta)\rho(x|s) \\
&\qquad - \sum_i u_i(\omega,\theta) \left(\sum_{s_i'} \alpha_i^*(s_i|s_i')\sigma(s_i', s_{-i}, \theta) - \sum_{s_i'} \alpha_i^*(s_i'|s_i)\sigma(s,\theta) \right) \rho(x|s) \Bigg] \\
&= \sum_s \max_\omega \sum_{\theta,x} \Bigg[w(\omega,\theta)\sigma(s,\theta)\rho(x|s) + \sum_i u_i(\omega,\theta)\beta_i^*(s_i)\sigma(s,\theta)\rho(x|s)
\end{aligned}
$$

[12] The latter infimum is across all ordered information structures and multipliers C.

$$- \sum_i u_i(\omega,\theta)\left(\sum_{s_i'} \alpha_i^*(s_i|s_i')\sigma(s_i', s_{-i},\theta) - \sum_{s_i'} \alpha_i^*(s_i'|s_i)\sigma(s,\theta)\right)\rho(x|s)\Bigg]$$

$$= P(S,\sigma).$$

The first equality comes from the fact that θ is uncorrelated with x given s, and the second equality comes from the optimality of α^* and β^*.

□

The infinite signal $x_i = \infty$ in the dual reduction represents all of the signals in (X,σ) that do not commute to a signal with a binding participation constraint. This represents a subtle difference with the ordered information structures and constraints used in the upper bounding program of Brooks and Du (2024), where it is assumed that all signals communicate with a binding participation constraint. Under this assumption, the two constructions coincide (up to a rescaling of the multiplier on local outward incentive constraints).

The construction of the informational virtual objective just presented is therefore more general. It seems that the infinite signal cannot be dispensed with, and it could be that all participation constraints are slack at the potential minimizer. Such an example is given in Section 4.3 of Brooks and Du (2024), in the context of a public goods problem, where the potential is minimized by an information structure in which each agent has a single signal, meaning there is no information about the state. Curiously, for this example, the duality gap is positive. We conjecture that there may be a more general connection between whether there is a duality gap and whether participation constraints are slack at the potential minimizer.

Finally, as with the mechanism, we can modify our construction to yield a truncated dual reduction that has a virtually lower potential than the original information structure. Given a dual reduction $(\overline{X}^*, \sigma^*)$, we define its *k-truncation* to be the ordered information structure $(\overline{X}^*, \widehat{\sigma})$, where

$$\widehat{\sigma}(x,\theta) = \sum_{y\in(\zeta^k)^{-1}(x)} \sigma^*(y,\theta);$$

$$\zeta_i^k(x) = \begin{cases} k & \text{if } k < x_i < \infty; \\ x_i & \text{otherwise.} \end{cases}$$

Note that for all k, the k-truncation is a finite information structure.

Proposition 2. *For any finite information structure $I = (S,\sigma)$ and corresponding dual reduction $(\overline{X}^*, \sigma^*)$ and for any $\epsilon > 0$, there exists a k so that if $(\overline{X}^*, \widehat{\sigma})$ is the k-truncation of $(\overline{X}^*, \sigma^*)$, then $\overline{P}(\overline{X}^*, \widehat{\sigma}) \leqslant P(I) + \epsilon$.*

As a result, the infimum potential across all finite information structures is equal to the infimum expected highest informational virtual objective across all finite ordered information structures.

Proof. Let $\overline{w} = \max_{\theta,\omega} w(\omega,\theta)$ and $\overline{u} = \max_{i,\theta,\omega} u_i(\omega,\theta)$. The expected highest informational virtual objective for the k-truncation $(\overline{X}^*, \widehat{\sigma})$ is

$$\begin{aligned}
\overline{P}(\overline{X}^*, \widehat{\sigma}) = \sum_x \max_\omega \sum_\theta \Bigg[& w(\omega,\theta)\widehat{\sigma}(x,\theta) + 2C\mathbb{I}_{x_i=0} \sum_i u_i(\omega,\theta)\widehat{\sigma}(x,\theta) \\
& + C\mathbb{I}_{x_i>0} \sum_i u_i(\omega,\theta) \left[\widehat{\sigma}(x,\theta) - \widehat{\sigma}(x_i - 1, x_{-i}, \theta)\right] \Bigg] \\
\leqslant \sum_x \max_\omega \sum_\theta \Bigg[& w(\omega,\theta)\sigma^*(x,\theta) + 2C\mathbb{I}_{x_i=0} \sum_i u_i(\omega,\theta)\sigma^*(x,\theta) \\
& + C\mathbb{I}_{x_i>0} \sum_i u_i(\omega,\theta) \left[\sigma^*(x,\theta) - \sigma^*(x_i - 1, x_{-i}, \theta)\right] \Bigg] \\
& + (\overline{w} + 4C\overline{u}) \sum_{x,\theta} |\sigma^*(x,\theta) - \widehat{\sigma}(x,\theta)| \\
\leqslant \overline{P}(\overline{X}^*, \sigma^*) + 2(\overline{w} & + 4C\overline{u}) \sum_\theta \sum_{\{x|k<x_i<\infty \text{ for some } i\}} \sigma^*(x,\theta).
\end{aligned}$$

Now, it must be that for k sufficiently large,

$$\sum_\theta \sum_{\{x|k<x_i<\infty \text{ for some } i\}} \sigma^*(x,\theta) < \frac{\epsilon}{2(\overline{w} + 4C\overline{u})},$$

since otherwise σ^* could not integrate to one. Thus, by Theorem 2, we have that

$$\overline{P}(\overline{X}^*, \widehat{\sigma}) \leqslant P(I) + \epsilon,$$

as desired. □

5 Discussion

Theorems 1 and 2 show that in solving for the maximum guarantee and minimum potential, it is without loss to assume that there is a simple one-dimensional structure on binding equilibrium constraints. This provides a foundation for the tools developed in Brooks and Du (2024), namely, that one can solve for the maximum guarantee by maximizing the expected lowest strategic virtual objective across all ordered participation secure mechanisms, and that one can solve for the minimum potential by minimizing the expected highest informational virtual objective across all ordered information structures.

While the bounds are always tight in this sense, it might still be the case that the minimum potential is strictly greater than the maximum guarantee. Indeed, Brooks and Du (2024) give an example where there is such a positive "duality gap." They also prove non-constructively that there is no duality gap for a wide class of optimal auctions problems. It remains an important and open direction for future research, to establish useful sufficient conditions for max guarantee to equal min potential.

More broadly, we have provided a deeper understanding of those environments that are especially challenging for a mechanism designer, namely, those in which types are linearly ordered by a "willingness" to participate in the mechanism. And we have provided a deeper understanding of those indirect which are informationally-robust in the presence of binding participation constraints, namely, those in which actions are linearly ordered by a "degree of participation" in the mechanism.

Finally, we comment on the connection with Myerson (1997).[13] That paper considers the correlated equilibria of a complete information normal-form game. Correlated equilibria are joint distributions over actions that satisfy obedience constraints. As in our analysis, Myerson interprets the multipliers on those obedience constraints as a Markov chain, and derives from it a particular "dual reduction" game, where actions in the reduced game correspond to mixtures in the original game. In fact, these mixtures are the invariant measures that are induced by the Markov chain. A main result of that paper is that any correlated equilibrium of the reduction is associated with a correlated equilibrium of the original game. In that sense, dual reduction shrinks the set of correlated equilibria.

It is natural to apply this idea to informationally-robust mechanism design, since a reduction in the set of BCE would necessarily be associated with an increase in the guarantee. Indeed, there is no great difficulty in adapting Myerson's construction to the setting where there is a payoff relevant state θ, and we consider BCE instead of correlated equilibria. The problem is that a reduction in Myerson's sense might no longer be participation secure. For example, consider the common value first-price auction, whose revenue guarantee was computed by Bergemann, Brooks, and Morris (2017). There are many ways in which we could "reduce" the first-price auction so as to shrink the set of BCE, e.g., by forcing all players to bid the ex ante expected value (which is indeed a BCE, induced by an equilibrium when the bidders have no information). But such a reduction would obviously fail to satisfy natural participation constraints when bidders do have information about the value.[14] One can view our dual reduction as a way of shrinking the set of feasible outcomes and improving the guarantee without losing participation security.[15] Indeed, as $x_i \to \infty$,

[13]More recently, Myerson (2024) extends that work to communication equilibria of sender-receiver games, and offers a "dual reduction" of the sender's information that is analogous to our dual reduction of an information structure.

[14]It is important to distinguish the optimal multipliers for the revenue guarantee program, which were used in our dual reduction, versus the multipliers used in the construction of the invariant measures in Myerson (1997). In general the two sets of multipliers are distinct. If we looked for invariant measures with respect to the optimal multipliers for the revenue guarantee, as identified by Bergemann, Brooks, and Morris (2017), the only invariant measure would be to bid the highest amount in the support of the revenue minimizing BCE, which is greater than the ex ante expected value and clearly not a BCE of the first-price auction.

[15]Clearly, by constraining the agents to only playing certain mixtures in the original game, we reduce the set of feasible joint distributions over actions and outcomes. We do not know whether this construction shrinks the set of equilibrium outcomes. The proof that the dual reduction in Myerson (1997) reduces the set of correlated equilibria relies on the fact that the reduced actions are invariant measures, and moreover, that the multipliers on obedience constraints induce $\lambda = 0$. But in general, $\lambda \neq 0$ for our optimal solution. But this issue is not relevant to our primary concern, which is achieving a higher guarantee. Similarly, we do not know whether our dual reduction of the information structure reduces the set of outcomes that can be implemented in equilibrium, but we do know that the best implementable outcome for the dual reduction is weakly worse than that for the original information structure.

the reduced action converges to the invariant measure for the recurrent class that is reached from the participation secure action in the original mechanism. In a similar spirit, our dual reduction of the information structure shrinks the set of feasible outcomes and reduces the best implementable objective for the designer, without weakening the agents' participation constraints.

References

Bergemann, D., B. Brooks, and S. Morris (2017): "First-Price Auctions with General Information Structures: Implications for Bidding and Revenue," *Econometrica*, 85, 107–143.

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

Brooks, B. and S. Du (2021a): "Maxmin Auction Design with Known Expected Values," Tech. rep., The University of Chicago and University of California-San Diego, working paper.

——— (2021b): "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.

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

Brooks, B., S. Du, and A. Haberman (2024a): "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 (2024b): "An Informationally Robust Model of Perfect Competition," Tech. rep., The University of Chicago and University of California-San Diego, working paper.

Bulow, J. and P. Klemperer (1996): "Auctions Versus Negotiations," *The American Economic Review*, 180–194.

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

Myerson, R. B. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, 6, 58–73.

——— (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.