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

Benjamin Brooks Songzi Du

December 9, 2025

Abstract

The *guarantee* of a mechanism is the lowest possible performance, across all information structures and equilibria. We define a *dual reduction* of a mechanism, which is itself a mechanism. We show that a dual reduction of mechanisms necessarily increases the guarantee; moreover, actions can be ordered so that the improvement in the guarantee is implied by the equilibrium conditions that whatever action an agent takes, they do not prefer to take the next higher action.

The *potential* of an information structure is the highest possible performance, across all mechanisms and equilibria. We also define a dual reduction of an information structure, which is itself an information structure. We show that dual reduction of information structures necessarily lowers the potential; signals are ordered, and the decrease in the potential is implied by agents not wanting to misreport the next lower signal and the interim participation constraint for the lowest signal.

Our results imply that in maximizing the guarantee and in minimizing the potential, it is without loss of generality to restrict attention to ordered mechanisms and ordered information structures that are *regular*, in the sense that the guarantee or the potential is determined by uni-directional local equilibrium constraints.

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. We have benefitted tremendously from conversations with Roger Myerson and Phil Reny. This research has been supported by NSF #2215259 and #2215475.

1 Introduction

In the standard approach to Bayesian mechanism design, there is a group of agents who have private information, and a mechanism designer who wishes to elicit that information and use it to improve their decision making. For many problems of interest, such as auctions or bilateral trade, it is also typical to suppose that the agents can object to the mechanism and obtain some outside option, so that the designer must also structure the mechanism to overcome the agents' objections. A key ingredient of the model is an *information structure*, which is comprehensive description of the agents' private information: the possible payoff-relevant states of the world, and the possible things that agents might know about the state and what others know. Given such a description, we can then proceed to analyze what outcomes can be attained under different mechanisms and equilibria. In particular, we can solve for those mechanisms and equilibria that achieve the optimal value for the designer, which we refer to as the information structure's *potential*.

In our view, any particular information structure and equilibrium are not to be taken too literally. Rather, these concepts operate together as a metaphor for forces towards common knowledge and mutual best responses that we think do exist in the world and would tend to shape behavior in long-lived institutions. And they guide our exploration of scenarios of how agents might behave under a given mechanism. With this perspective, it seems appropriate to understand how a given mechanism might perform under various information structures and equilibria. A recent literature on informational robustness has embraced this perspective in a rather extreme form, with a focus on understanding a mechanism's lowest possible performance, across all information structures and equilibria, which has been termed the mechanism's *guarantee*. And whereas in the classical approach we would optimize the performance under a particular information structure and equilibria (the latter selected by the designer), the informationally-robust approach instead seeks to identify those mechanisms that achieve the highest possible guarantee. Such guarantee-maximizing mechanisms necessarily have structure that drives their performance that is not tied to the details of any particular information structure and equilibrium.

In this paper, we investigate the structure of mechanisms that maximize the guarantee. Also, from a complementary perspective, we also consider those information structures that minimize the potential. This dual pair of optimization problems allows us to approach from both sides the issue of what can be achieved by a mechanism designer who is unwilling or unable to commit to a single description of information and equilibrium. Our setting is quite general: There is an abstract set of states, an abstract set of outcomes, and an expected utility preference for each the agents and for the designer.

We have two main results, one for guarantees and one for potentials. For the former, suppose that we are given a mechanism and its guarantee. As we review below, the guarantee is the solution to a linear program, so associated with the optimal solution to that program are optimal Lagrange multipliers on equilibrium constraints. We use the Lagrange multipliers to define a new mechanism, which we term the *dual reduction*. In fact, this is just the original mechanism, but where the agents' are restricted to play one of a carefully curated *ordered sequence* of mixed strategies. We prove that the dual reduction necessarily has a weakly higher guarantee than the original mechanism, while still affording the agents' the same capacity to object to the mechanism if they so wish. But what is especially interesting is the manner in which we verify that the guarantee is higher. In the dual reduction, the agents' actions are ordered, such that lowest action registers the strongest objection to the mechanism—what we refer to as a *secure* action, analogous to bidding zero in an auction. And to prove that the dual reduction's guarantee is higher than that of the original mechanism, we only rely on the fact that in any equilibrium, when an agent is taking a given action, they are not made better of by a "local" deviation to the action one step farther away from the secure action.

This result has a powerful implication: Consider now the problem of maximizing the guarantee over all mechanisms. For any mechanism we may consider, it has a dual reduction which attains an even higher guarantee. The dual reduction's actions are ordered, with the lowest action being secure, and the improvement in the guarantee is certified by a first-order lower bound that is derived from local equilibrium constraints. This tells us that the maximum guarantee must be *equal to* the maximum first-order lower bound across all ordered mechanisms. Finally, and remarkably, we show that maximizing the first-order lower bound across ordered mechanisms can be reduced to solving a family of finite-dimensional linear programming problems, parametrized by a number of actions and a (constant) Lagrange multiplier on the local equilibrium constraints. As the number of actions and the Lagrange multiplier go to infinity, the values of these linear programs converge to the supremum guarantee, and the sequence of solutions to the linear programs provide approximate guarantee-maximizing mechanisms.

We now describe the dual result for information. Suppose we are given an information structure and its potential. Again, the potential is the optimal value of a linear program. Using the associated optimal Lagrange multipliers on equilibrium and participation constraints, we define a new dual-reduction information structure. In the dual reduction, each agent's signal is a noisy observation of their signal in the original information structure. Moreover, the reduced signals are ordered in a manner that captures, in a certain sense, the degree to which the agent's preferences are opposed to those of the designer. We prove that the dual reduction necessarily has a weakly lower potential than the original information structure. And to prove that the dual reduction's potential is lower, we only rely on the fact that in any equilibrium, when an agent observes the signal associated with the strongest opposition, they are willing to participate in the mechanism, and otherwise, they are not made better off by mimicking the behavior when the signal is one step closer to that with the strongest opposition. Thus, at the minimum potential, the signal with strongest opposition is made indifferent to leaving the mechanism; the next signal strictly prefers to participate in the mechanism, but is indifferent to acting *as if* they are indifferent to leaving the mechanism; the next signal strictly prefers participation and is indifferent to acting as if they are indifferent to acting as if they are indifferent to leaving the mechanism; and so on.

Once again, there is a broader implication of this result: Consider the problem of minimizing the potential over all information structures. For any information structure we may consider, it has a dual reduction which attains an even lower potential. The dual reduction's signals are ordered, and the reduction in the potential is certified by a first-order upper bound that is derived from local equilibrium constraints constraints, and the requirement that the agent with the lowest signal is willing to participate. Finally, we show that minimizing the first-order order upper bound on the potential across ordered mechanisms can be reduced to solving a one-dimensional family of finite-dimensional linear programming problems, parametrized by a number of signals and a (constant) Lagrange multiplier on the participation and local equilibrium constraints. As the number of signals and the multiplier go to infinity, the values of these linear programs converge to the infimum potential, and the sequence of solutions to the linear programs provide approximate potential-minimizing information structures.

The emphasis on local equilibrium constraints evokes various results in classical Bayesian mechanism design, where under certain assumptions on preferences and private information, we can conclude that local equilibrium constraints are sufficient to pin down the designer's optimal payoff. The most famous such result is that of Myerson (1981), in the context of optimal auctions for a single unit of a good with independent private values: in the so-called "regular" case, optimal revenue is equal to a first-order upper bound on the potential. In our model, both the mechanism and the information structure are endogenous objects. However, our result on dual reductions of information structures shows the following: there are always potential-minimizing information structures that are *regular*[1] in that signals can

[1] In our formal treatment below, the supremum potential is not attained exactly. Both potential maximization and regularity hold approximately. A similar qualification applies to our statements regarding guarantee maximization and regularity of mechanisms.

be ordered and the potential is pinned down by a first-order upper bound. In fact, for a general class of revenue-maximization problems with interdependent values and transferable utility, we show that the potential-minimizing information will always feature independently distributed signals, and the potential is equal to an expected highest virtual value, in the sense of Myerson (1981). Beyond the optimal auctions problem, we show that the minimum potential is equal to a minimum expected highest *informational virtual objective*, which generalizes the virtual value of Myerson (1981), Bulow and Klemperer (1996), and others.

Analogously, there also exist guarantee-maximizing mechanisms that are regular, in that the optimal guarantee is pinned down from a first-order lower bound on the potential. Moreover, there is a counterpart to the informational virtual objective, which is the *strategic virtual objective*.

Both the strategic and informational virtual objectives appeared in our prior work, where we analyzed various informationally robust mechanism design problems. For example, in Brooks and Du (2021), we solved for guarantee-maximizing mechanisms and potential-minimizing information structures for common value auctions. In Brooks and Du (2024), we described a general methodology for solving such problems, which were precisely the first-order upper and lower bounds on the potential and guarantee, respectively. However, Brooks and Du (2024) did not prove that these bounds were tight in general.[2] The present paper uses dual reductions to prove that these bounding programs are *always* tight. This provides a solid foundation for first-order approach in informationally-robust optimal mechanism design.

The rest of this paper is structured as follows. Section 2 introduces notation and terminology. Section 3 presents our result on dual reductions of mechanisms.

In Section 4, we discuss a distinct notion of dual reduction of Myerson (1997), which inspired our analysis.

Section 5 presents our analysis of dual reductions of information structures.

In Section 6, we offer an interpretation of our results is demonstrating that max guarantees and min potentials will select mechanisms and information structures that are regular, in a sense generalizing that which is well-known in the literature.

In Section 7, we generalize our results to the case where there is either an upper bound or a lower bound on the agents' information. The upper bound takes the form of an individual garbling complete set of information structures, as described by Brooks, Du, and Haberman (2024a); in this case, our results go through essentially unchanged. For the lower bound, we suppose that the agents observe signals from a fixed baseline information

[2] Brooks and Du (2024) showed that the bounding programs were tight for a class of optimal auctions problems.

structure, and may observe more, as in Bergemann and Morris (2016). The lower bound significantly enriches our theory: We provide analogous results concerning dual reductions, but the agents' actions or signals now take the form of *sequences* of signals in the base information structure, and we discuss implications and simulations for revenue-guarantee maximizing auctions with private values. In particular, we present simulations in which the max guarantee is strictly higher than the highest revenue achievable with a dominant strategy mechanism.

Section 8 is a conclusion, and an appendix contains omitted proofs.

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 prior distribution of θ is denoted $\mu \in \Delta(\Theta)$ and is held fixed throughout our analysis.[3]

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 $S = \prod_i S_i$,[4] 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.[5] We further assume that every $s_i \in S_i$ has positive probability, i.e., $\sum_{s_{-i},\theta} \sigma(s_i, s_{-i}, \theta) > 0$. 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.[6]

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

[4] 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, $i = 1, \ldots, N$.

[5] 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.

[6] 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.

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* (or just *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 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 finite mechanisms is $\overline{\mathcal{M}}$, 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 strategy profile b is *interim individually rational* if $U_i(M, I, b) \geqslant 0$ for all i. We let the set of interim individually rational equilibria be denoted by $\mathcal{E}^{IR}(M, I)$.

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 the highest payoff that the designer can achieve across all mechanisms and equilibria subject to the agents being willing to participate. In the present theory, "willing to participate" will be operationalized as an interim individual rationality constraint.

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

The convention is that if $\mathcal{E}^{IR}$ is empty, then the inner supremum is $-\infty$. Note that the potential is necessarily greater than the designer's maximum payoff across all participation secure mechanisms and equilibria.[7] This is because if M is participation secure, then $\mathcal{E}(M, I) = \mathcal{E}^{IR}(M, I)$ for every I, simply because the option to deviate to a participation secure action implies that $U_i(M, I, b) \geqslant 0$. It follows immediately that for any $M \in \overline{\mathcal{M}}^*$ and $I \in \overline{\mathcal{I}}$, $G(M) \leqslant P(I)$.

As is well-known, given a mechanism $M = (A, m)$, we can apply the revelation principle for information design to the computation of $G(M)$: When minimizing the designer's welfare over information structures and equilibria, it is without loss to consider *direct recommendation* information structures for which the space of signals is $S = A$ and the *obedient* strategies for which $b_i(a|a) = 1$ are an equilibrium (Bergemann and Morris, 2016). In particular, $G(M)$ for $M = (A, m)$ is the solution to the following (infinite dimensional) linear programming problem. Let $\Delta_\mu(A)$ be the set of distributions on $A \times \Theta$ with marginal μ. Then

$$\begin{aligned} G(M) = \inf_{\sigma \in \Delta_\mu(A)} & \sum_{\theta, a, \omega} w(\omega, \theta) m(\omega|a) \sigma(a, \theta) \\ \text{s.t.} & \sum_{\theta, a_{-i}, \omega} u_i(\omega, \theta) \left[m(\omega|a_i, a_{-i}) - m(\omega|a_i', a_{-i}) \right] \sigma(a_i, a_{-i}, \theta) \geqslant 0 \ \forall i, a_i, a_i'. \end{aligned}$$

The linear inequalities are referred to as *obedience constraints*. Note that the optimization problem is an infimum because the mechanism may have countably many actions, so that there are countably infinitely many variables and constraints. However, if the mechanism is finite, then this is a finite dimensional linear program, and the infimum is a minimum.

Analogously, applying the revelation principle for mechanism design (Myerson, 1981), $P(I)$ for $I = (S, \sigma)$ can be computed by optimizing over *direct revelation* mechanisms for which $A = S$ and the *truthful* strategies for which $b_i(s_i|s_i) = 1$ are an equilibrium. Thus,

[7] In prior work, we have *defined* the potential to be the designer's maximum payoff across all participation secure mechanisms and equilibria. However, in operationalizing the potential, we have always relaxed participation security to an interim participation constraint, and for the cases where there is zero duality gap, the infimum potential is the same for both definitions. To define dual reductions of information structures, it is essential that we use the definition in terms of interim participation.

$P(I)$ is the solution to the following linear programming problem:

$$
\begin{aligned}
P(I) = & \sup_{m:S\to\Delta(\Omega)} \sum_{\theta,s,\omega} w(\omega,\theta)m(\omega|s)\sigma(s,\theta) \\
\text{s.t.} & \sum_{\theta,s_{-i},\omega} u_i(\omega,\theta)\left[m(\omega|s_i,s_{-i}) - m(\omega|s_i',s_{-i})\right]\sigma(s_i,s_{-i},\theta) \geqslant 0 \ \forall i,s_i,s_i' \\
& \quad \sum_{\theta,s_{-i},\omega} u_i(\omega,\theta)m(\omega|s_i,s_{-i})\sigma(s_i,s_{-i},\theta) \geqslant 0 \ \forall i,s_i.
\end{aligned}
$$

The truthtelling constraints are in the second line and interim participation constraints are the third line. Again, there is a supremum because S may be infinite. If S is finite, then this supremum is a maximum.

3 Dual Reductions of Mechanisms

This section presents our first main result: For any mechanism M, we can construct a *dual reduction* of it which necessarily has a higher guarantee.

3.1 The Dual Reduction

We first define a dual reduction. Given a finite mechanism $M = (A,m)$, 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,

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

Let

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

and without loss, set

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

Thus, we have chosen $\alpha_i^*(a_i|a_i)$ so that for every a_i, the multipliers on deviations from a_i have the same sum, regardless of a_i. We refer to C as the *deviation flow* associated with M and the dual reduction. We note that C is a kind of "gross" flow, because it includes within it flow $\alpha_i(a_i|a_i)$, which is a flow that actually results in not deviating, whereas $C-1$ represents a "net" deviation flow that corresponds to deviations to $a_i' \neq a_i$.

Let $X_i = \{0, 1, 2, \ldots\}$ to be the non-negative integers, and $X = \prod_i X_i$. The dual reduction will have an action space equal to X. For each i, let a_i^0 be (any) participation secure action for player i. Then, for each i, define a function $b_i^* : X_i \to \Delta(A_i)$ as follows: $b_i^*(a_i^0|0) = 1$, and for $k > 0$,

$$b_i^*(a_i|k) = \sum_{a_i'} b_i^*(a_i'|k-1)\frac{\alpha_i^*(a_i|a_i')}{C}.$$

Finally, we define for all ω and x,

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

where

$$b^*(a|x) = \prod_i b_i^*(a_i|x_i).$$

This completes the definition of a dual reduction (X, m^*).

The dual reduction can be interpreted in the following way. We can view $\alpha_i^*(\cdot|a_i)/C$ as a particular stochastic deviation, wherein whenever the action a_i would have been played, the agent deviates to a_i' with probability $\alpha_i^*(a_i'|a_i)/C$. This stochastic deviation is the one that is "most tempting" for agent i, in the sense that the optimal Lagrangian for $G(M)$ can be obtained by attaching a Lagrange multiplier of C to this stochastic deviation, and setting all other obedience multipliers to zero. Now, in the dual reduction, the action $x_i = 0$ corresponds to taking a participation secure action a_i^0 with probability 1; $x_i = 1$ corresponds to starting from a_i^0, and drawing a new action from the most tempting deviation from a_i^0; inductively, action x_i corresponds to drawing an action from mixture associated with $x_i - 1$ and then taking the most tempting deviation from action drawn. Thus, actions in the dual reduction are interpreted as a number of iterated most tempting deviations, starting from a_i^0.

We note that there is more than one dual reduction because there may be many choices of saddle points for the guarantee program, and in particular, it is always possible to make

C (and the $\alpha_i^*(s_i|s_i)$ multipliers) larger, and there may be more than one choice for the secure action that we place at the beginning of the sequence.

3.2 Ordered Mechanisms and the Strategic Virtual Objective

The dual reduction is an example of what we refer to as an *ordered mechanism*: a mechanism for which the action space is X, and 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.

Our analysis of dual reductions of mechanisms will rely on a first-order lower bound on the guarantee for ordered mechanisms that was developed in Brooks and Du (2024). Fix an ordered mechanism (X, m) and a constant $C \in \mathbb{R}_+$. For a given $C > 0$, the associated *strategic virtual objective* at the action profile x and the state θ is

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

This object 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.

We note that the notation C for the Lagrange multiplier on local constraints in the strategic virtual objective is intentionally the same notation that we used for the deviation flow in the dual reduction. In the proof of Theorem 1 below, they will turn out to be the same quantity.[8]

We have the following lemma:[9]

[8]The constant C is effectively the multiplier on the agents' local 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. Indeed, the construction of the dual reduction 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

[9]Note in the case where (X, m) is not finite, we have not ruled out the possibility that there *are no BCE*, in which case $G(X, m) = \infty$. This is not a logical problem for Theorem 1 below, although it does raise a concern that results concerning guarantees over equilibrium outcomes may be vacuous. As we shall see, Proposition 1 shows that our key insights go through, with all due approximations, if we restrict attention to finite mechanisms, so that such concerns are ultimately moot.

Lemma 1. *For any ordered mechanism (X, m) and $C > 0$,*

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

This result was proved in Proposition 3 of Brooks and Du (2024) for the case of finite ordered mechanisms. The proof is short:

Proof. For any BCE of (X, m), i, and a_i, we have that

$$\sum_{x_{-i}, \omega, \theta} \sigma(x_i, x_{-i}, \theta) u_i(\omega, \theta)(m(\omega|x_i + 1, x_{-i}) - m(\omega|x_i, x_{-i})) \leqslant 0.$$

Hence, the designer's payoff is at least[10]

$$\begin{aligned}
&\sum_{x,\theta,\omega} \sigma(x, \theta) \left[w(\omega, \theta) m(\omega|x) + C \sum_{i} u_i(\omega, \theta)(m(\omega|x_i + 1, x_{-i}) - m(\omega|x)) \right] \\
&= \sum_{x,\theta} \sigma(x, \theta) \lambda(x, \theta, C) \geqslant \sum_{\theta} \mu(\theta) \inf_{x} \lambda(x, \theta, C),
\end{aligned}$$

as desired. □

Thus, the guarantee of an ordered mechanism is always at least the expected lowest strategic virtual objective. It is $\underline{G}(X, m, C)$ that we have referred to in the introduction as the first-order lower bound.

3.3 Main Result

We can now state our main result regarding dual reductions of mechanisms:

Theorem 1. *Fix a finite participation secure mechanism $M = (A, m)$. For any dual reduction (X, m^*) and associated deviation flow C, we have that*

$$G(X, m^*) \geqslant \underline{G}(X, m^*, C) \geqslant G(M).$$

Proof. By Lemma 1, it suffices to prove the second inequality. We have:

$\underline{G}(X, m^*)$

[10] 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.

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

$$= \sum_\theta \mu(\theta) \inf_x \sum_{\omega,a} \left[w(\omega,\theta) m(\omega|a) b^*(a|x) + C \sum_i m(\omega|a) \left(b_i^*(a_i|x_i+1) - b_i^*(a_i|x_i) \right) b_{-i}^*(a_{-i}|x_{-i}) u_i(\omega,\theta) \right]$$

$$= \sum_\theta \mu(\theta) \inf_x \sum_{\omega,a} \left[w(\omega,\theta) m(\omega|a) b^*(a|x) \right.$$

$$\left. + C \sum_i m(\omega|a) \left(\sum_{a_i'} \frac{\alpha_i^*(a_i|a_i')}{C} b_i^*(a_i'|x_i) - b_i^*(a_i|x_i) \right) b_{-i}^*(a_{-i}|x_{-i}) u_i(\omega,\theta) \right]$$

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

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 proof shows that dual reduction leads to a higher guarantee through two channels. First, dual reduction reduces the set of outcomes which are feasible, since the agents are only allowed to use certain mixtures in the original mechanism. Second, these mixtures are chosen so that the critical stochastic deviation α^*/C is still feasible for the agents. These two properties together imply that minimum welfare in the dual reduction is higher than that in the original mechanism.

3.4 Finite Approximations

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 $M = (A, m)$, with dual reduction (X, m^*) and associated deviation flow C, 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 b^* for $(X, m^* \circ f^k)$ is $\widehat{b}_i^*(a_i|x_i) = b_i^*(a_i|f^k(x_i))$. Let the associated mechanism be denoted by $\widehat{m} = m^* \circ f^k$:

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

Let C be the deviation flow associated with (X, m^*). We compare the strategic virtual objectives of (X, m^*) and $(X, \widehat{m})$, which are:

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

and

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

$$|\lambda^*(x, \theta, C) - \widehat{\lambda}(x, \theta, C)|$$

$$\leqslant M \sum_{\omega,a} |b^*(a|x) - \widehat{b}^*(a|x)| + CM \sum_{\omega,a,i} (|b_i^*(a_i|x_i+1) - b_i^*(a_i|x_i)|\mathbb{I}_{x_i \geqslant k} + |b_{-i}^*(a_{-i}|x_{-i}) - \widehat{b}_{-i}^*(a_{-i}|x_{-i})|),$$

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 chains (e.g., Stroock, 2014, equation (4.1.15) on page 85), $\lim_{k\to\infty} b_i^*(k)$ 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 $b_i^*(\infty)$.

Since

$$|b_i^*(a_i|x_i) - \widehat{b}_i^*(a_i|x_i)| \leqslant |b_i^*(a_i|x_i) - \widehat{b}_i^*(a_i|k)|\mathbb{I}_{x_i \geqslant k}$$

and

$$\lim_{k\to\infty} b_i^*(a_i|k) = b_i^*(a_i|\infty)$$

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. □

3.5 Computing Approximate Guarantee Maximizers

Taken together, Theorem 1 and Proposition 1 imply that we can compute the supremum guarantee via a sequence of finite-dimensional linear programs parametrized by the deviation flow C and number of actions k. Let us define $X_i(k) = \{0,\ldots,k-1\}$ and $X(k) = \prod_i X_i(k)$. The linear program is:

$$\max_{m:X(k)\to\Delta(\Omega),\lambda:\Theta\to\mathbb{R}} \sum_\theta \lambda(\theta)\mu(\theta)$$

$$\text{s.t. } \lambda(\theta) \leqslant \sum_\omega \Big[w(\omega,\theta)m(\omega|x) + C\sum_i (m(\omega|\min\{x_i+1,k-1\},x_{-i}) - m(\omega|x))u_i(\omega,\theta) \Big] \ \forall\theta, x. \tag{1}$$

Let us denote by $G^*(k,C)$ the value of this linear program. Then we have the following result:

Corollary 1. *For all ϵ, there exists a $\underline{k}$ and $\underline{C}$ so that if $k \geqslant \underline{k}$ and $C \geqslant \underline{C}$, then $G^*(k,C)+\epsilon$ is greater than the supremum guarantee. Moreover, if m solves the linear program* (1) *for* (k,C)*, then $G(X(k),m) \geqslant G^*(k,C)$.*

3.6 An Example: First-Price Auctions

As an illustration, we may consider the revenue guarantee of the first-price auction with discrete bids in a finite grid between 0 and $\overline{a} > 0$. The state θ is the vector of bidders' ex post values for the good $(\theta_1, \ldots, \theta_N)$, and the outcome ω consists of vectors of allocation probabilities $q = (q_1, \ldots, q_N) \in \mathbb{R}_+$ with $\sum_i q_i = 1$ and payments $t = (t_1, \ldots, t_N)$. And given a bid profile a, the mechanism m puts equal probability on the outcomes where one of the high bidders i is allocated the good and pays their bid $t_i = a_i$, and all other bidders pay nothing. Revenue is $w(\omega, \theta) = \sum_i t_i$ and bidder i's utility is $u_i(\omega, \theta) = \theta_i q_i - t_i$. Note that the bid $a_i = 0$ is a secure action for each player.

For the case where the prior μ on θ is exchangeable, Bergemann, Brooks, and Morris (2017) computed the revenue guarantee of the first-price auction and the associated revenue-minimizing BCE. While they analyzed a model with continuous bids, we may heuristically apply their results to the discrete case. A key finding of Bergemann, Brooks, and Morris (2017) is that in the BCE that attains the revenue guarantee, $\alpha_i(a_i'|a_i) > 0$ for $a_i \neq a_i'$ if and only if $a_i < a_i' \leqslant \widehat{a}$, where $\widehat{a}$ is the highest bid in the support of the equilibrium. Extrapolating to the discrete case, a dual reduction of the first-price auction would consist of restricting the bidders to a sequence of mixtures over bids. The first element of the sequence is probability one on bidding zero. The remaining elements of the sequence are full support mixtures over bids that are increasing in the first-order stochastic dominance order. In the limit, this sequence of mixtures converges weakly to a mixture that puts probability one on bidding $\widehat{a}$.

3.7 Coarse Guarantees

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 2. *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 Comparison with the Dual Reduction of Myerson (1997)

Myerson (1997) considers the correlated equilibria of a complete information normal-form game.[11] 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 his dual reduction, these mixtures are actually the invariant measures under the Markov chain. A key finding in that paper is that any correlated equilibrium of the reduction is effectively a correlated equilibrium of the original game, in the sense that the players would be willing to randomize conditionally independently, according to the mixtures they were recommended in the correlated equilibrium of the dual reduction. 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

[11] 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.

problem is that a reduction in Myerson's sense might no longer be participation secure. For example, let us reconsider the common value first-price auction. 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.[12] One can view our dual reduction as a way of shrinking the set of feasible outcomes in a manner that improves the guarantee, without compromising on participation security.[13] In a similar spirit, our dual reduction of the information structure, developed in the next section, shrinks the set of feasible outcomes that can be implemented and reduces the best equilibrium outcome for the designer, without weakening the agents' participation constraints.

5 Dual Reductions of Information Structures

We now present an analogous theory of dual reductions of information structures.

5.1 The Dual Reduction

We first 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 optimal multipliers $\alpha_i^*(s_i'|s_i)$ and $\beta_i^*(s_i)$ such that

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

[12]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, these 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.

[13]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.

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

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

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

Let $\overline{X}_i = \{0, 1, 2, \ldots\} \cup \{\infty\}$. In other words, we take X and add a point at infinity. 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,

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

As before, we refer to C as a *deviation flow.* Starting from s_i, with probability $\beta_i^*(s_i)/C$, 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 consisting only of signals s_i' for which $\beta_i(s_i') = 0$.

We can define ρ recursively as follows: $\rho_i(0|s_i) = \beta_i^*(s_i)/C$, and for $x_i > 0$, it is given by

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

Finally, $\rho_i(\infty|s_i) = 1 - \sum_{x_i=0}^{\infty} \rho_i(x_i|s_i)$.

We then define

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

where

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

The dual reduction is $(\overline{X}, \sigma^*)$. In 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 took to transition to $\varnothing$ (or they observe ∞ if the signal never transitions to $\varnothing$).

We once again note that there is more than one dual reduction of an information structure because there is more than one saddle point of the potential linear program (in particular there is no bound on how large $\alpha_i^*(s_i|s_i)$ and C can be).

5.2 Ordered Information Structures and the Informational Virtual Objective

The dual reduction is an example of an *ordered information structure* of the form $(\overline{X}, \sigma)$. Such an information structure is *finite* if there exists a k such that $\sigma(x,\theta) > 0$ only if $x_i < k$ for all i.

Our analysis of dual reductions of information structures will rely on a first-order upper bound on the potential for ordered information structures that was developed in Brooks and Du (2024). 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

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

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.[14]

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

We have the following Lemma, which is a trivial generalization of Proposition 3 of Brooks and Du (2024):

Lemma 2. *For any ordered information structure* $(\overline{X}, \sigma)$,

$$P(\overline{X}, \sigma) \leqslant \overline{P}(\overline{X}, \sigma, C) \equiv \sum_x \max_\omega \psi(x, \omega, C).$$

Proof. For any incentive compatible and individually rational direct mechanism on $(\overline{X}, \sigma)$, we must have for all i

$$\sum_{x_{-i}} \sigma(0, x_{-i}, \theta) u_i(\omega, \theta) m(\omega | 0, x_{-i}) \geqslant 0$$

and for all i and $x_i > 0$,

$$\sum_{x_{-i}} \sigma(x_i, x_{-i}, \theta) u_i(\omega, \theta)(m(\omega | x_i, x_{-i}) - m(\omega | x_i - 1, x_{-i})) \geqslant 0.$$

Hence, the designer's payoff is at most

$$\begin{aligned}
&\sum_{x,\theta,\omega} \sigma(x, \theta) \Bigg[w(\omega, \theta) m(\omega | x) + C \sum_i \mathbb{I}_{x_i = 0} u_i(\omega, \theta) m(\omega | x) \\
&\qquad\qquad + C \sum_i \mathbb{I}_{x_i > 0} u_i(\omega, \theta)(m(\omega | x) - m(\omega | x_i - 1, x_{-i})) \Bigg] \\
&= \sum_{x,\theta,\omega} m(\omega | x) \Bigg[w(\omega, \theta) \sigma(x, \theta) + C \sum_i u_i(\omega, \theta)(\sigma(x_i + 1, x_{-i}, \theta) - \sigma(x, \theta)) \Bigg] \\
&\leqslant \sum_x \max_\omega \sum_\theta \Bigg[w(\omega, \theta) \sigma(x, \theta) + C \sum_i u_i(\omega, \theta)(\sigma(x_i + 1, x_{-i}, \theta) - \sigma(x, \theta)) \Bigg] \\
&= \sum_x \max_\omega \psi(x, \omega, C),
\end{aligned}$$

as desired. □

Thus, the potential of an ordered information structure is always at most the expected highest informational virtual objective.[15]

[15] As with Lemma 1, we note that Lemma 2 does not rule out the possibility that there are no incentive compatible and individually rational direct mechanisms, in which case we would have $P(I) = -\infty$. The concern that Theorem 2 might be vacuous will be similarly dispelled by Proposition 2, which shows that our results hold with arbitrarily small error when we restrict attention to finite approximations of dual reductions.

5.3 Main Result

We can now state our main result regarding dual reductions of information structures:

Theorem 2. *Fix a finite information structure $I = (S, \sigma)$. For any dual reduction $(\overline{X}, \sigma^*)$ and associated deviation flow C, we have that $P(\overline{X}, \sigma^*) \leqslant \overline{P}(\overline{X}, \sigma^*, C) \leqslant P(I)$.*

Proof. By Lemma 2, it suffices to prove the second inequality. For the reduced information structure, we now have the following upper bound on designer welfare:

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

The above equation is clearly weakly less than

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

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

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

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

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 $x_i = \infty$, for which no constraints bind, was not part of the description of an ordered mechanism and the associated first-order upper bound in Brooks and Du (2024). The construction of the informational virtual objective presented in this paper is therefore more general. The infinite signal arises naturally from dual reductions, whenever there are signals in the original information structure that do not commute with a signal that has a binding participation constraint. Moreover, it seems that the infinite signal cannot be dispensed with. In fact, 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. For this example, the potential is minimized by an information structure in which each agent has a single signal, meaning there is no information about the state, and all incentive and participation constraints are slack.

5.4 Finite Approximations

As with mechanisms, 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 $(\overline{X}, \sigma^*)$ with deviation flow C, 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 $\overline{P}(\overline{X}, \widehat{\sigma}, C) \leqslant P(I) + \epsilon$.*

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 $\overline{u} = \max_{i,\theta,\omega} u_i(\omega, \theta)$. Letting C be the associated deviation flow, the expected highest informational virtual objective for the k-truncation $(\overline{X}, \widehat{\sigma})$ is

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

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. □

5.5 Computing Approximate Potential Minimizers

Taken together, Theorem 2 and Proposition 2 imply that we can compute the infimum potential via a sequence of finite-dimensional linear programs parametrized by the deviation flow C and a number of signals k. Let us define $\overline{X}_i(k) = \{0, \ldots, k-2, \infty\}$ and $\overline{X}(k) = \prod_i \overline{X}_i(k)$. The linear program is:

$$\begin{aligned} &\min_{\sigma \in \Delta_\mu(\overline{X}(k)), \psi : \overline{X}(k) \to \mathbb{R}} \sum_x \psi(x) \\ \text{s.t. } &\psi(x) \geqslant \sum_\theta \Big[w(\omega, \theta)\sigma(x, \theta) - C \sum_i \mathbb{I}_{x_i < \infty} u_i(\omega, \theta)(\sigma(x_i + 1, x_{-i}, \theta) - \sigma(x, \theta) \Big] \; \forall \omega, x, \end{aligned} \tag{2}$$

with the convention that $\sigma(k-1, x_{-i}, \theta) = 0$. Let us denote by $P^*(k, C)$ the value of this linear program. Then we have the following result:

Corollary 3. *For all ϵ, there exists a $\underline{k}$ and $\underline{C}$ so that if $k \geqslant \underline{k}$ and $C \geqslant \underline{C}$, then $P^*(k, C) - \epsilon$ is less than the infimum potential. Moreover, if σ solves the linear program (2) for (k, C), then $P(\overline{X}(k), \sigma) \leqslant P^*(k, C)$.*

5.6 An Example: Optimal Auctions

To illustrate dual reductions of information, we will return to the optimal auctions problem, in which a single unit of a good is for sale. As in Section 3.6, the state is a vector of values for the good, and the outcome consists of a vector of allocation probabilities and transfers. The designer's objective is to maximize revenue, and the agents have the standard quasilinear utility. In our baseline model, there are finitely many outcomes ω, but following the standard convention in optimal auctions, we will regard the transfers as free variables, so that the potential is

$$\begin{aligned} P(I) = &\max_{q : S \to \Delta(\{0,1,\ldots,N\}), t : S \to \mathbb{R}^N} \sum_{s,\theta} \sigma(s, \theta) \Bigg[\sum_i t_i(s) + \beta_i^*(s_i)(q_i(s)\theta_i - t_i(s)) \\ &+ \sum_{i, s_i'} \alpha_i^*(s_i' | s_i) \left[(q_i(s) - q_i(s_i', s_{-i}))\theta_i - (t_i(s) - t_i(s_i', s_{-i})) \right] \Bigg] \\ = &\sum_s \max_{q \in \Delta(\{0,1,\ldots,N\}), t \in \mathbb{R}^N} \sum_\theta \Bigg[\left(\sum_i t_i \right) \sigma(s, \theta) + \sum_i \beta_i^*(s_i)(q_i\theta_i - t_i)\sigma(s, \theta) \\ &+ \sum_{i, s_i'} \left[\alpha_i^*(s_i' | s_i)\sigma(s, \theta) - \alpha_i^*(s_i | s_i')\sigma(s_i', s_{-i}, \theta) \right] (q_i\theta_i - t_i) \Bigg], \end{aligned}$$

for the optimal multipliers α^* and β^*. Note that for every finite information structure, a saddle point exists for the potential linear program because, first, there exists a feasible solution (the degenerate mechanism in which the good is not allocated and no one pays is both incentive compatible and individually rational) and second, in any individually rational mechanism, revenue is bounded above by the expectation of the highest θ_i.

Brooks and Du (2024) generalize the first-order upper bound on the potential to the optimal auctions problem.[16] The informational virtual objective at an outcome $\omega = (q, t)$ reduces to

$$\psi(x, q, t, C) = \sum_{\theta} \left(\sum_{i} t_i \right) \sigma(x, \theta) - C \sum_{i} (q_i \theta_i - t_i)(\sigma(x_i + 1, x_{-i}, \theta) - \sigma(x, \theta)).$$

Let us denote the marginal distribution on x by

$$\eta(x) = \sum_{\theta} \sigma(x, \theta).$$

Then, the coefficient on t_i in the informational virtual objective at x is

$$\eta(x) - C(\eta(x_i + 1, x_{-i}) - \eta(x)). \tag{3}$$

The maximum informational virtual objective at (x, θ) is finite only if the coefficients on t_i are all zero, meaning that η solves the difference equation that (3) is equal to zero at all x. Otherwise, if the coefficient is non-zero, it would be possible to drive the Lagrangian to infinity by either sending t_i to ∞ (if the coefficient is positive) or $-\infty$ (if the coefficient is negative). In fact, this difference equation has a solution if and only if $C > 1$, in which case the unique solution is that the x_i are distributed independent geometric with arrival rate $1/C$, and there is zero probability on the infinite type $x_i = \infty$.[17] Explicitly,

$$\eta(x) = C^{-N} \left(\frac{C - 1}{C} \right)^{\sum_i x_i}.$$

How can we reconcile this result with dual reductions? Fix an information structure I and an associated dual reduction I^*. Theorem 7 in Appendix B is an analogue of

[16] A superficial difference between the present model and Brooks and Du (2024) is that the latter take C to be a positive integer and label the signals in increments of $1/C$.

[17] If $x_i = \infty$, then we have $\eta(x) - C(\eta(x_i + 1, x_{-i}) - \eta(x)) = \eta(x)$ (as $x_i - 1 = x_i$). Note that in Brooks and Du (2024), the number of signals is finite, in which case the marginal over signals that causes the transfer to drop out of the Lagrangian is a censored geometric, where all of the mass on signals bigger than the highest signal is pooled on the highest signal.

Theorem 2 for the optimal auctions problem with unbounded transfers. It shows that $P(I^*) \leqslant \overline{P}(I^*, C) \leqslant P(I)$. So, in particular, the upper bound on the potential for I^* is finite, and hence, $C > 1$ and η must be iid geometric! Why would this be the case? Well, if $C < 1$, then for every s_i, we would have $\sum_{s_i' \neq s_i} \alpha_i^*(s_i'|s_i) + \beta_i^*(s_i) < 1$. Thus, the coefficient on $t_i(s_i)$ in the Lagrangian for the revenue maximization problem would be

$$1 - \sum_{s_i'} \alpha_i^*(s_i'|s_i) - \beta_i^*(s_i) > 0,$$

so that the optimal value of the Lagrangian would be infinite. This contradicts the hypothesis that (α^*, β^*) are part of a saddle point. Hence, the optimal multipliers must have the property that $\sum_{s_i' \neq s_i} \alpha_i^*(s_i'|s_i) + \beta_i^*(s_i) = 1$, so that $C = 1 + \alpha_i^*(s_i|s_i) > 1$.

We can also give a direct argument for why, in dual reductions for the optimal auctions problem, the infinite signals must have zero likelihood. To see this, note that the infinite signal can have positive probability only if the Markov chain b_i^* has a recurrent class $\widehat{S}_i \subseteq S_i$ for which $\beta_i^*(s_i) = 0$ for all $s_i \in \widehat{S}_i$. Signals in $\widehat{S}_i$ would be mapped onto the infinite type in the dual reduction.[18] Now, suppose that such a recurrent class exists, and let $f_i(s_i)$ be an invariant measure under b_i^* supported on $\widehat{S}_i$. Then

$$f_i(s_i) = \sum_{s_i'} b_i^*(s_i|s_i') f_i(s_i'),$$

and hence (using the fact that $\beta_i^*(s_i) = 0$ for all s_i in the support of f_i)

$$\sum_{s_i'} (\alpha_i^*(s_i|s_i') - \alpha_i^*(s_i'|s_i)) f_i(s_i) = 0.$$

Now, consider a mechanism in which no agent is ever allocated the good, and the only transfer is that player i pays $\kappa f_i(s_i)$ for some $\kappa > 0$. Since the all signals in $\widehat{S}_i$ have positive probability, revenue is positive, but $\beta_i^*(s_i) = 0$ for all s_i in $\widehat{S}_i$, so that the value of the Lagrangian for $P(I)$ is

$$\kappa \sum_{s,\theta} \sigma(s,\theta) \left[f_i(s_i) - \sum_{s_i'} \alpha_i^*(s_i'|s_i)(f_i(s_i) - f_i(s_i')) \right] = \kappa \sum_{s,\theta} \sigma(s,\theta) f_i(s_i).$$

[18]It could also be that $\rho_i(\infty|s_i) > 0$ for s_i that are not in such a recurrent class, if there is positive probability of transitioning from such an s_i to a recurrent class with no binding participation constraints.

Thus, by taking $\kappa \to \infty$, we could drive the value of the Lagrangian to infinity, which contradicts finiteness of the revenue potential.

Independence of the signals in the dual reduction is more subtle. While it follows from our theorem, we can shed light on it with an example. Suppose we start with an information structure I where the signals are highly correlated. In particular, let us go so far as to suppose that the information structure satisfies the hypotheses of Crémer and McLean (1988) for full-surplus extraction.[19] In that case, it is well known that incentive constraints are slack at the optimal mechanism,[20] so that $\alpha_i^*(s_i'|s_i) = 0$ for all i, s_i, and $s_i' \neq s_i$. So the only way for transfers to drop out of the Lagrangian for the revenue maximization problem would be for $\beta_i^*(s_i) = 1$ for all i and s_i, and these are indeed the optimal multipliers. In our construction of the dual reduction, we would then pick $C > 1$, and set $\alpha_i^*(s_i|s_i) = C - 1$. Thus, no matter which signal s_i is drawn initially, exit occurs with probability $1/C$ in each round, and hence the exit time is independent of the initial signal.

Incidentally, this argument shows that from an information structure with full-surplus extraction, the dual reduction will be one in which the agents have no information at all about θ. The resulting optimal revenue is necessarily lower, because total surplus is minimized if the agents know nothing about θ (and hence it is impossible to target the allocation towards agents with higher values).

More generally, we can conclude that in any dual reduction, the signals will be iid geometric with arrival rate $1/C$, so that transfers drop out of the informational virtual objective. For x finite, let us write

$$v_i(x) = \frac{1}{\eta(x)} \sum_\theta \sigma(x, \theta)\theta_i$$

for the interim expected value. Then the remaining terms in the informational virtual objective are

$$\begin{aligned} \psi(x, q, t, C) &= C \sum_i q_i \left[v_i(x)\eta(x) - v_i(x_i + 1, x_{-i})\eta(x_i + 1, x_{-i}) \right] \\ &= \sum_i q_i \left[v_i(x)\eta(x) + v_i(x)\eta(x)(C - 1) - C v_i(x_i + 1, x_{-i})\eta(x_i + 1, x_{-i}) \right] \end{aligned}$$

[19] In particular, each signal s_i's beliefs about s_{-i} are not a convex combination of the beliefs held by agent i when they observe other signals.

[20] Indeed, Crémer and McLean (1988) construct transfers so that agents strictly prefer reporting their true signal to any misreport.

$$= \sum_i q_i \left[v_i(x) - \frac{C-1}{C} C \left(v_i(x_i+1, x_{-i}) - v_i(x) \right) \right] \eta(x).$$

This expression is a discrete analogue of the virtual value of Myerson (1981). Indeed, Bulow and Klemperer (1996) generalize the virtual value to the case where agents have real-valued signals that are independently distributed according to an absolutely continuous cumulative distribution $F_i(x_i)$, and there is a differentiable value function. They show that the virtual value has the form

$$v_i(x) - \frac{1-F_i(x_i)}{f_i(x_i)} \frac{\partial v_i(x)}{\partial x_i}, \tag{4}$$

where f_i is the probability density function corresponding to F_i. In the special case where $x_i = v_i(x)$, this reduces to the standard private value formula in Myerson (1981). Comparing (4) with the final expression for ψ, we can relate the two formulae if $dx \approx 1/C$, i.e., $1/C$ is our scaling of one increment in the signal. In particular, the (discrete) inverse hazard rate of the geometric distribution with arrival rate is $1/C$, so that $(C-1)/C$ is analogous to $(1 - F_i(x_i))/f_i(x_i)$.

Our point is that dual reductions of information structures—inspired by the dual reductions of Myerson (1997)—inevitably lead us to the conclusion that for the optimal auctions problem, the minimum potential is attained with independent signals, and the minimum potential itself is an expected highest virtual value of the buyer who is allocated the good—a structure that echos and generalizes the celebrated revenue equivalence formula of Myerson (1981) for independent private values.

6 A Regular Interpretation

In the classic auction design problem with independent private values, signals are equal to ex post values, and therefore are ordered as real numbers. A key step in the analysis of Myerson (1981) is that local downward equilibrium constraints can be used to solve out transfers in terms of the allocation rule, and reformulate the expected revenue of a direct mechanism as the expected *virtual value* of the agent who is allocated the good. Thus, one can get quite far in the analysis working with only the local downward equilibrium constraints. But in general, the local downward constraints are not sufficient to pin down optimal revenue. However, if the value distribution is log concave, then local downward constraints are sufficient, and optimal revenue is simply the expected highest virtual value. Following Myerson (1981), this has come to be known as the regular case.

As we showed in the previous section, the informational virtual objective is a natural generalization of the virtual value to general mechanism design problems: it is the designer's payoff from the induced outcome, plus additional terms corresponding to local "downward" deviations, towards a "lowest" type that has a binding participation constraint. As is the case in optimal auctions, for any given ordered information structure, the designer's optimal payoff is generally *not* equal to the expected highest informational virtual objective. However, an implication of Theorem 2 is that this is always the case for potential-minimizing information structures, in the following approximate sense.

For $\epsilon > 0$, we say that an information structure I is an *ϵ-potential minimizer* if $P(I) \leqslant \inf_{I' \text{ finite}} P(I') + \epsilon$. We say that I is *ϵ-regular* if for some $C > 0$, $\overline{P}(I, C) \leqslant P(I) + \epsilon$.

Corollary 4. *For every $\epsilon > 0$, there exists a finite information structure that is an ϵ-potential minimizer and is ϵ-regular.*

Proof. Let I be an $\epsilon/2$-potential minimizer, and let I^* be a dual reduction of I with deviation flow C. By Theorem 2, we have $P(I^*) \leqslant \overline{P}(I^*, C) \leqslant P(I) \leqslant \inf_{I' \text{ finite}} P(I') + \epsilon/2$. By Proposition 2, there exists a k such that if $\widehat{I}$ is the k-truncation of I^*, then $P(\widehat{I}) \leqslant \overline{P}(\widehat{I}, C) \leqslant P(I) + \epsilon/2$. Hence,

$$\inf_{I' \text{ finite}} P(I') \leqslant P(\widehat{I}) \leqslant \overline{P}(\widehat{I}, C) \leqslant \inf_{I' \text{ finite}} P(I') + \epsilon,$$

so that $\widehat{I}$ is a finite ϵ-potential minimizer that is also ϵ-regular. □

There is of course an analogous result for mechanisms. For $\epsilon > 0$, a mechanism M is an *ϵ-guarantee maximizer* if $G(M) \geqslant \sup_{M' \text{ finite, participation secure}} G(M') - \epsilon$. And we say that M is *ϵ-regular* if $G(M) \geqslant \underline{G}(M, C) - \epsilon$. By an analogous argument, we have the following:

Corollary 5. *For every $\epsilon > 0$, there exists a finite mechanism that is an ϵ-guarantee maximizer and is ϵ-regular.*

Regularity is typically conceived of as a condition on primitives that *implies* an orderly structure on equilibrium constraints, and a variety of such conditions have been proposed in the literature.[21] In the definitions and corollaries just presented, we are offering a new perspective on the issue: Rather than looking for conditions on primitives directly, we may instead *define* regularity as the property that the binding equilibrium constraints are local and uni-directional. Our theorems show that such regular structure will naturally and

[21]See, for example, Myerson (1981), Bulow and Klemperer (1996), Chung and Ely (2007), Yamashita and Zhu (2018), Chen and Li (2018), Bergemann et al. (2018), Yang (2023), and Loertscher and Muir (2024).

endogenously arise in information structures that minimize the potential and in mechanisms that maximize the guarantee. This is true regardless of what assumptions one makes on the primitives Θ, Ω, u_i, and w.

7 Bounds on Information

The guarantee is the minimum value to the designer across all information structures and equilibria. In applications of the theory, it may be desirable to consider the worst case over only a subset of information structures, which will be regarded as *admissible*. Following Brooks, Du, and Haberman (2024a), we may distinguish between lower bounds on information (i.e., assuming that agents always have some minimal amount of information) and upper bounds on information (i.e., assuming that agents may be bounded away from full information about the state). We will consider each of these in turn.

7.1 Upper Bounds on Information

For upper bounds on information, we suppose that the admissible set of information structures is *individual garbling complete*, in the sense described in Brooks, Du, and Haberman (2024a). This property essentially says that if an information structure I is admissible, and if I' represents the agents observing conditionally independent noisy signals about their information in I (an *individual garbling*), then there is an admissible information structure that is equivalent to I'.[22]

7.1.1 Guarantees with Upper Bounds

Any set of admissible information structures induces a *feasibility correspondence* that associates to each action space A a set of outcomes $F(A) \subseteq \Delta(A \times \Theta)$ that can be induced by some admissible information structure and some strategies of the players. Theorem 1 of Brooks, Du, and Haberman (2024a) shows that the admissible set of information structures is individual garbling complete if and only if, for every game, the set of equilibrium outcomes across all admissible information structures is equal to the set of BCE that are also feasible. Moreover, Theorem 2 of that paper shows a feasibility correspondence is induced by a set of information structures that is individual garbling complete if and only if F is itself individual garbling complete, in the following sense: for any $\sigma \in \Delta(A \times \Theta)$, and for

[22]Two information structures are equivalent if each is an individual garblings of the other.

any strategies $b_i : A_i \to \Delta(A'_i)$, the outcome defined by $\sigma'(a', \theta) = \sum_a b(a'|a)\sigma(a, \theta)$ is in $F(A')$.

Thus, in analyzing guarantees with information upper bounds of this form, we may model the restriction on information implicitly, by taking as given an individual garbling complete feasibility correspondence F. We further suppose that F is convex valued which, as shown in Brooks, Du, and Haberman (2024a), is equivalent to assuming that the admissible set of information structures being *public randomization complete*. We further suppose that for each A, $F(A)$ is *finitely generated*, meaning that it is the intersection of finitely many linear constraints. The additional hypotheses of convexity and finite generation imply that for a finite mechanism, the guarantee program with feasibility constraints is still a finite-dimensional linear programming problem. An example of a correspondence that satisfies these hypotheses would be the correspondence of outcomes with bounded total variation distance with the product of marginals, as studied in Brooks, Du, and Haberman (2024a).

The restricted guarantee of a mechanism $M = (A, m)$ is $G_F(M)$, which is the minimum value of the designer across all BCE that are in $F(A)$.[23] This is again the solution to a linear programming problem. We also define, for an ordered mechanism (X, m), the lower bound

$$\underline{G}_F(X, m, C) = \inf_{\sigma \in F(X)} \sum_{\theta, x} \sigma(x, \theta)\lambda(x, \theta, C).$$

By walking through the same steps as for Lemma 1, we conclude that this is indeed a lower bound on $G_F(X, m)$.

Now, given a mechanism M, define an F-dual reduction to be a dual reduction with respect to the optimal Lagrange multipliers on obedience constraints in the linear program for $G_F(M)$. We have the following theorem:

Theorem 3. *Let F be an individual garbling complete, convex-valued, and finitely generated feasibility correspondence. For any finite participation secure mechanism $M = (A, m)$ and corresponding deviation flow C and dual reduction (X, m^*), we have that $G_F(X, m^*) \geqslant \underline{G}_F(X, m^*, C) \geqslant G_F(M)$.*

Proof. Exactly the same steps as in the proof of Theorem 1 show that the Lagrangian for $\underline{G}_F(X, m^*, C)$ is the same as that for $G_F(M)$. For each i, let b_i^* be the mapping constructed

[23] If F is individual garbling complete, then F includes all outcomes where a and θ are independent. Thus, $F(A)$ will contain BCE that are correlated equilibria of the game where the agents have no information about the state.

in that proof from X_i to mixtures in $\Delta(A_i)$. Then for any feasible outcome σ^* of (X, m^*), there is another feasible outcome σ of M, which is the pushforward of σ^* through the mixtures $b^*(x)$. This outcome attains the same value in the Lagrangian for M as σ^* does in the lower bound Lagrangian for M^*. Finally, σ is an individual garbling of σ^*, with the mixtures $b^*(x)$ playing the role of the garbling. Therefore, σ is feasible for the program $G_F(M)$, so that $G_F(M) \leqslant \underline{G}_F(X, m^*, C)$. □

We also have an analogue of Proposition 1, that we can approach $G_F(M)$ with truncations of (X, m^*). As a corollary of Theorem 3, we know that even with the upper bound on information, the supremum of $G_F(M)$ over all finite participation secure M is equal to the supremum of $\underline{G}_F(M, C)$ over all finite ordered M and $C > 0$. One difference, however, is that with the upper bound on information, the outcome that pointwise minimizes the strategic virtual objective, state by state, may no longer be feasible. Thus, it is no longer exactly correct to say that the guarantee supremizers are those that supremize the expected lowest strategic virtual objective. It is however correct to say that they supremize the lowest expected strategic virtual objective, where lowest is evaluated across all feasible outcomes.[24]

7.1.2 Potentials with Upper Bounds on Information

We now explore the implications of upper bounds on information for potentials. An admissible set of information structures is *closed under individual garbling* if for any admissible I and individual garbling I' of I, I' is also admissible. Closed under individual garbling is more demanding than individual garbling completeness, as defined in Brooks, Du, and Haberman (2024a). However, this is the right notion for characterizing potential-minimizing information structures.[25,26] It is also without loss, in that Proposition 1 of

[24] To operationalize the upper bound on information, we would have to incorporate the feasibility constraints into the Lagrangian. This could be done by adding extra Lagrange multipliers to the strategic virtual objective, corresponding to those constraints. See Brooks, Du, and Haberman (2024a) for examples.

[25] A set of information structures is individual garbling complete if for every admissible I and individual garbling I' of I, there is an admissible I'' such that I' is a *coordinated individual garbling* of I''. Roughly speaking, a coordinated individual garbling is one that is self-reinforcing, in that each agent regards the information being garbled away as useless, provided that the other agents' signals are being garbled. This definition is the right one for characterizing the set of all equilibrium outcomes that can be attained with information in the set: if I' is a coordinated individual garbling of I'', then every equilibrium outcome that can be induced under I' can also be induced under I''. However, this does not mean that the two information structures have the same potential. Indeed, it could be that I' is a coordinated individual garbling of I'' and $P(I') < P(I'')$. Thus, for the purpose of calculating infimum potentials, it is important to work with a stronger notion of completeness, that the admissible set of information structures is closed under individual garbling.

[26] We could weaken the condition by only requiring I' to be *equivalent* to some admissible I'', in the sense defined in Brooks, Du, and Haberman (2024a) and Lehrer, Rosenberg, and Shmaya (2013), but this

Brooks, Du, and Haberman (2024a) shows that any feasibility correspondence F that is individual garbling complete is induced by a set of information structures that is closed under individual garbling (and in particular, F is induced by the admissible set of direct recommendation information structures, corresponding to outcomes in the range of F).

The last point we wish to make is that for any information structure I and associated dual reduction I^*, I^* is in fact an individual garbling of I. Hence, if $I = (S, \sigma) \in F(S)$ and if F is individual garbling complete, then $I^* = (X^*, \sigma^*) \in F(X^*)$. Thus, for any admissible information structure, there is an ordered information structure that is also admissible and achieves a lower potential. We conclude that the infimum $P(I)$ across all finite admissible information structures is equal to the infimum of $\overline{P}(I', C)$ across all admissible finite ordered information structures I' and $C > 0$. We summarize this discussion with the following result:

Theorem 4. *Suppose that $\mathcal{I}$ is a set of information structures that is closed under individual garbling. Fix a finite information structure $I = (S, \sigma) \in \mathcal{I}$. For any deviation flow C and corresponding dual reduction $(\overline{X}, \sigma^*)$, we have that $P(\overline{X}, \sigma^*) \leqslant \overline{P}(\overline{X}, \sigma^*, C) \leqslant P(I)$ and $(\overline{X}, \sigma^*) \in \mathcal{I}$.*

7.2 Lower Bounds on Information

We now consider a lower bound on information. In particular, following Bergemann and Morris (2016), we suppose that there is a base information structure $I = (S, \sigma)$. The agents observe their signals from I but may also have more information. Thus, the set of admissible information structures are those of the form $I' = (\widehat{S}, \sigma')$, where $\widehat{S} \subseteq S \times S'$ for some s', and the marginal of σ' on $S \times \Theta$ is σ. This is a lower bound in the *individual sufficiency* order. In contrast to the upper bounds on information, lower bounds on information will substantially enrich our theory.

7.2.1 Guarantees with Lower Bounds

We now denote by $G_I(M)$ the guarantee when only information structures that are above I in the in the individual sufficiency order are admissible. As shown by Bergemann and Morris (2016), for a finite mechanism M, and if S is finite, $G_I(M)$ is the minimum welfare across all BCE where the obedience constraints are conditioned on both a_i and s_i. This is a finite dimensional linear program that is feasible and bounded, so it necessarily has a

would not affect the infimum potential; it would only affect the particular choice of labels for signals in the potential infimizing information structures.

saddle point. In particular, there are Lagrange multipliers $\alpha_i^* : A_i \times A_i \times S_i \to \mathbb{R}_+$ such that

$$G_I(M) \equiv \min_{\substack{\sigma' \in \Delta(A\times S\times\Theta) \\ \text{s.t. } \operatorname{marg}_{S\times\Theta} \sigma' = \sigma}} \sum_{a,s,\theta,\omega} \sigma'(a,s,\theta) \Bigg[w(\omega,\theta) m(\omega|a) + \sum_{i,a_i'} \alpha_i^*(a_i'|a_i,s_i) u_i(\omega,\theta)(m(\omega|a_i',a_{-i}) - m(\omega|a_i,a_{-i})) \Bigg]$$

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

Dual reductions now must have a richer space of actions. In particular, actions are finite sequences of base types: $X_i = \cup_{k=1}^\infty S_i^k$. A representative k-length sequence will be denoted $\mathbf{s}_i = (s_i^0, s_i^1, \ldots, s_i^{k-1})$. Each finite sequence is associated with a lottery $b_i^*(\mathbf{s}_i) \in \Delta(A_i)$. As before, fix an action a_i^0 that is participation secure. For each s_i, $b_i^*(\cdot|s_i)$ assigns probability one to a_i^0, and for other sequences, we define b_i^* recursively:

$$b_i^*(a_i|\mathbf{s}_i, s_i) = \sum_{a_i'} \alpha_i^*(a_i|a_i', s_i) b_i^*(a_i'|\mathbf{s}_i).$$

For an action profile a and a sequence profile $\mathbf{s}$, we define $b^*(a|\mathbf{s}) = \prod_i b_i^*(a_i|\mathbf{s}_i)$. Finally, the dual reduction's outcome mapping is

$$m^*(\omega|\mathbf{s}) = \sum_a m(\omega|a) b^*(a|\mathbf{s}).$$

Now, we claim that the restricted guarantee $G_I(X, m^*)$ is at least $G_I(M)$. As before, we prove this using a lower bound $\underline{G}_I(M^*, C)$, where we drop all obedience constraints except those associated with a deviation from the action $\mathbf{s}_i$ to $(\mathbf{s}_i, s_i)$ (i.e., adding another element to the sequence) and attaching a multiplier C to those constraints. Explicitly, the lower bound is

$$\underline{G}_I(X, m^*, C) = \sum_{s,\theta} \sigma(s,\theta) \inf_{\mathbf{s}} \sum_\omega \left[w(\omega,\theta) m^*(\omega|\mathbf{s}) + C \sum_i u_i(\omega,\theta) \left(m^*(\omega|(\mathbf{s}_i, s_i), \mathbf{s}_{-i}) - m^*(\omega|\mathbf{s}) \right) \right].$$

Theorem 5. *Fix a lower bound information structure I. For any finite participation secure mechanism $M = (A, m)$ and corresponding deviation flow C and dual reduction (X, m^*), $G_I(X, m^*) \geqslant \underline{G}_I(X, m^*, C) \geqslant G_I(M)$.*

The proof follows that of Theorem 1 and appears in the Appendix.

Why have we defined the dual reduction in this manner? The high level concept behind the dual reduction is that we restrict the mechanism by constraining the agents to play only a subset of mixed strategies, but we do so in such a manner that the "most tempting deviations" α_i^*/C are still feasible for the agents. That is enough to ensure that the guarantee in the dual reduction must weakly increase. But in the model with base signals, each s_i has a different most tempting deviation. So, if we include the pure mixture a_i^0, then we must also include $\alpha_i^*(\cdot|a_i^0, s_i)$ for every i. Such mixtures correspond to the sequences of length 2, of the form (s_i^0, s_i). But now, one possibility is that in equilibrium agent i plays $\alpha_i^*(\cdot|a_i^0, s_i)$, but their true base signal is s_i' and they are most tempted to deviate according to $\alpha_i^*(\cdot|\cdot, s_i')$. This would correspond to a sequence of the form (s_i^0, s_i, s_i'), and so on. Thus, the sequences correspond to iterated most tempting deviations of a particular sequence of base signals.

Notice that if $|S_i| = 1$ for all i, then the lower bound is degenerate, and we are back to the baseline model where the lower bound is uninformative. Indeed, in that case, the actions in the dual reduction would be sequences of the form $(s_i, \ldots, s_i)$, for the single signal $s_i \in S_i$. Such sequences could equivalently be labeled according to their length, just as we did in Section 3.

7.2.2 Potentials with Lower Bounds

For the potential-minimization problem, we now write out the optimal Lagrangian for the potential of an information structure that is individually sufficient for I:

$$
\begin{aligned}
P(I') = &\max_{m: S\times S' \to \Delta(\Omega)} \sum_{s,s',\theta,\omega} \sigma'(s,s',\theta)\Bigg[w(\omega,\theta)m(\omega|s,s') + \sum_i \beta_i^*(s_i,s_i')u_i(\omega,\theta)m(\omega|s,s') \\
&\quad + \sum_{i,(\widehat{s}_i,\widehat{s}_i')} \alpha_i^*(\widehat{s}_i,\widehat{s}_i'|s_i,s_i')u_i(\omega,\theta)(m(\omega|s,s') - m(\omega|\widehat{s}_i,s_{-i},\widehat{s}_i',s_{-i}')) \Bigg] \\
= &\sum_{s,s'} \max_\omega \sum_\theta \Bigg[w(\omega,\theta)\sigma'(s,s',\theta) + \sum_i \beta_i^*(s_i,s_i')u_i(\omega,\theta)\sigma'(s,s',\theta) \\
&\quad + \sum_{i,\widehat{s}_i,\widehat{s}_i'} \left[\alpha_i^*(\widehat{s}_i,\widehat{s}_i'|s_i,s_i')\sigma'(s,s',\theta) - \alpha_i^*(s_i,s_i'|\widehat{s}_i,\widehat{s}_i')\sigma'(\widehat{s}_i,s_{-i},\widehat{s}_i',s_{-i}'\theta)\right] u_i(\omega,\theta) \Bigg].
\end{aligned}
$$

We define $\alpha_i^*(s_i, s_i'|s_i, s_i') > 0$ so that for all i, s_i, s_i', there is a constant deviation flow:

$$C = \beta_i^*(s_i, s_i') + \sum_{\widehat{s}_i, \widehat{s}_i'} \alpha_i^*(\widehat{s}_i, \widehat{s}_i'|s_i, s_i').$$

We then define Markov chain as follows: First draw (s, s', θ) from σ'. Then run the chain and transition from (s_i, s_i') to $(\widehat{s}_i, \widehat{s}_i')$ with probability

$$b_i^*(\widehat{s}_i, \widehat{s}_i'|s_i, s_i') = \alpha_i^*(\widehat{s}_i, \widehat{s}_i'|s_i, s_i')/C$$

and exit the mechanism with probability

$$b_i^*(\varnothing|s_i, s_i') = \beta_i^*(s_i, s_i')/C.$$

In the dual reduction, each agent is then informed of the sequence $\mathbf{s}_i = (s_i^0, s_i^1, \ldots, s_i^k)$ of base signals that they transitioned to before they exited, where k is finite or infinite. We write $\overline{X}_i$ for the set of such sequences. The first element of the sequence is always the agent's true base signal that they started with.

Let us write $\rho_i(s_i^0, s_i^1, \ldots, s_i^k|s_i, s_i')$ for the likelihood of observing the finite sequence $(s_i^0, \ldots, s_i^k)$ conditional on the original types being (s_i, s_i'). This is defined recursively as

$$\rho_i(s_i^0|s_i, s_i') = \begin{cases} b_i^*(\varnothing|s_i, s_i') & \text{if } s_i = s_i^0; \\ 0 & \text{otherwise.} \end{cases}$$

and

$$\rho_i(s_i^0, \ldots, s_i^k|s_i^0, s_i') = \sum_{\widehat{s}_i'} \rho_i(s_i^1, \ldots, s_i^k|s_i^1, \widehat{s}_i') b_i^*(s_i^1, \widehat{s}_i'|s_i^0, s_i')$$

when k is finite. When k is infinite, we define a measure over infinite sequences in S_i^∞ according to the Markov chain, which we denote by $\rho(d\mathbf{s}|s)$. For the rest of our analysis, we focus on the case where every signal commutes with a signal s_i that has $\beta_i^*(s_i) > 0$, so that the infinite types do not exist. This will always be the case for the optimal auctions problem that we explore at the end of this section.

In the dual reduction $I^* = (\prod_i(\cup_{k=0}^\infty S_i^k), \sigma^*)$, the likelihood of a signal profile is

$$\sigma^*(\mathbf{s}, \theta) = \sum_{s'} \prod_i \rho_i(\mathbf{s}_i|\mathbf{s}_i^0, s_i') \sigma'(\mathbf{s}^0, s', \theta)$$

where $\mathbf{s}^0$ is the vector of first elements in the sequence.

We now look at a Lagrangian where the participation constraint binds with multiplier C for the types $\mathbf{s}_i$ of length 1, and for any type $\mathbf{s}_i = (s_i^0, \ldots, s_i^k)$ of length greater than 1, the only truthtelling constraint that binds is the one that points to $(s_i^1, \ldots, s_i^k)$, again with multiplier C:

$$\overline{P}(\overline{X}, \sigma^*, C) \equiv \sum_{\mathbf{s}} \max_{\omega} \sum_{\theta} \left[w(\omega,\theta)\sigma^*(\mathbf{s},\theta) + C \sum_i u_i(\omega,\theta) \left(\sigma^*(\mathbf{s},\theta) - \sum_{s_i} \sigma^*((s_i, \mathbf{s}_i), \mathbf{s}_{-i}, \theta) \right) \right].$$

Theorem 6. *Fix a lower bound information structure I, and let I' be an information structure that is individually sufficient for I. Suppose further that for the optimal multipliers (α^*, β^*), the associated Markov chain is such that all signals exit in finite time with probability one. Then for any dual reduction I^* and associated deviation flow C, we have that $P(\overline{X}, \sigma^*) \leqslant \overline{P}(\overline{X}, \sigma^*, C) \leqslant P(I')$.*

The proof closely mirros that of Theorem 2, and is in the Appendix.

Thus, with the lower bound on information, the dual reduction must be enriched to capture a particular sequence of most tempting deviations, which are now indexed by the base signals. In the dual reduction, agents always learn which base signal they drew, but then they are informed of a particular sequence of misreports of the base signal that they would have taken, after which they exited the mechanism.

In formulating Theorem 6, we have ruled out the case where the Markov chain never exits. This is with loss of generality, and it has the potential to add considerable complexity to the theory. Although we see no conceptual difficulty in formulating the dual reduction with infinite sequences of base signals, there are benchmark cases where it will not arise, as in the optimal auctions problem that we return to next.

7.2.3 An Example: Optimal Auctions with Private Values

As an application of dual reductions with lower bounds on information, we now consider the special case of the optimal auctions problem, where $\omega = (q, t_1, \ldots, t_N)$ consists of an allocation and transfers; $w(\omega,\theta) = \sum_i t_i$; and $u_i(\omega,\theta) = v_i(q,\omega) - t_i$. We assume the transfers are free variables and there are finitely many possible allocations. At this point, we can be agnostic about the precise details of the space in which the allocations live.

We will start with dual reductions of information structures. As in Section 5.6, we claim that for any information structure $I' = (\widehat{S}, \sigma')$ and saddle point, when we construct the Markov chain b_i^*, there will not be any recurrent classes of signals for which $\beta_i^*(s_i, s_i') = 0$ for every (s_i, s_i') in the class. We now sketch the proof of the claim, which is similar to the one we provided before: Suppose there were such a class. Then the chain b_i^* restricted to that class has an invariant measure, which we may call $f(s_i, s_i')$. Taking $f(s_i, s_i')$ to be zero outside of the recurrent class, we would have that for all $(\widehat{s}_i, \widehat{s}_i')$,

$$\left(\sum_{s_i, s_i'} \alpha(\widehat{s}_i, \widehat{s}_i' | s_i, s_i') \right) f(\widehat{s}_i, \widehat{s}_i') = \sum_{s_i, s_i'} \alpha(\widehat{s}_i, \widehat{s}_i' | s_i, s_i') f(s_i, s_i'). \tag{5}$$

Now, consider the optimal Lagrangian for $P(I')$:

$$\sum_{s,s',\theta} \sigma'(s, s', \theta) \Bigg[\sum_i t_i(s, s') + \sum_i \beta_i^*(s_i, s_i') u_i(\omega, \theta) m(\omega | s, s')$$
$$+ \sum_{i, (\widehat{s}_i, \widehat{s}_i')} \alpha_i^*(\widehat{s}_i, \widehat{s}_i' | s_i, s_i') \left[v_i(q, \theta)(m(q|s, s') - m(q|\widehat{s}_i, s_{-i}, \widehat{s}_i', s_{-i}')) - t_i(s, s') + t_i(\widehat{s}_i, s_{-i}, \widehat{s}_i', s_{-i}') \right] \Bigg].$$

Suppose we replace the transfer rule $t_i(s, s')$ with $t_i(s, s') + \kappa f_i(s_i, s_i')$. This clearly increases revenue. It does not affect the term for participation constraints, however, since $\beta_i^*(s_i, s_i') = 0$ for all (s_i, s_i') in the recurrent class which contains the support of f_i. Moreover, by (5), there is no net change to the terms for truthtelling constraints either. This proves the claim.

Thus, for the optimal auctions problem, there are no recurrent classes in which $\beta_i^* = 0$, so that in any dual reduction, the infinite sequences of base signals have zero probability.

In addition, the coefficient on the transfer $t_i(\mathbf{s})$ will be

$$\sum_\theta \left[\sigma^*(\mathbf{s}, \theta) - C \sum_i \left(\sigma^*(\mathbf{s}, \theta) - \sum_{s_i} \sigma^*((s_i, \mathbf{s}_i), \mathbf{s}_{-i}, \theta) \right) \right].$$

This coefficient must be zero for the transfer to drop out and for the optimal value of the Lagrangian to be finite. Equivalently, if we write $\eta^*(\mathbf{s}) = \sum_\theta \sigma^*(\mathbf{s}, \theta)$, then for every $\mathbf{s}$ and i we have

$$\sum_{s_i} \eta^*((s_i, \mathbf{s}_i), \mathbf{s}_{-i}) = \frac{C-1}{C} \eta^*(\mathbf{s}),$$

where again we need $C > 1$. In other words, the distribution of the *length* of the sequence is geometric with arrival rate $1/C$, and this is also the distribution of the continuation sequence conditional on any partial sequence. We also have the constraint that for all

(s, θ).

$$\sigma(s,\theta) = \sum_{\mathbf{s}} \sigma^*((s,\mathbf{s}),\theta).$$

Let us now further specialize to single unit auctions with private values, meaning that $v_i(q,\theta) = q_i\theta_i$, where $q_i \in \{0,1\}$ and $\sum_i q_i \leqslant 1$, and $\theta = s = \mathbf{s}^0$. Then, for fixed C, the problem of minimizing the potential upper bound for optimal auctions reduces to:

$$\begin{aligned} &\min_{\eta(\mathbf{s})\geqslant 0,\psi(\mathbf{s})\geqslant 0} \sum_{\mathbf{s}} \psi(\mathbf{s}) \\ \text{s.t. } &\sigma(s) = \sum_{\mathbf{s}} \eta(s,\mathbf{s}) \ \forall s \\ &\sum_{s_i} \eta((s_i,\mathbf{s}_i),\mathbf{s}_{-i}) = \frac{C-1}{C}\eta(\mathbf{s}) \ \forall \mathbf{s}, i \\ &\psi(\mathbf{s}) \geqslant C\left[\mathbf{s}_i^0\eta(\mathbf{s}) - \sum_{s_i} s_i\eta((s_i,\mathbf{s}_i),\mathbf{s}_{-i})\right] \ \forall s \end{aligned} \tag{6}$$

By Theorem 6, as we take C to infinity, this sequence of upper bounds must converge to the infimum potential in the optimal private value auctions problem.

Dual reductions of mechanisms are more straightforward. We will do this only for the case of the single unit private-value auction problem, for which $v_i(q,\theta) = q_i\theta_i$. Plugging in functional forms for preferences, the problem of maximizing the guarantee upper bound is

$$\begin{aligned} &\max_{q_i(\mathbf{s})\geqslant 0, t_i(\mathbf{s}),\lambda(s)} \sum_{s} \lambda(s)\sigma(s) \\ \text{s.t. } &\lambda(s) \leqslant \sum_i t_i(\mathbf{s}) \\ &\qquad + C\sum_i \Big(s_i[q_i((\mathbf{s}_i,s_i),\mathbf{s}_{-i}) - q_i(\mathbf{s})] - [t_i((\mathbf{s}_i,s_i),\mathbf{s}_{-i}) - t_i(\mathbf{s})]\Big) \ \forall s,\mathbf{s} \\ &\sum_i q_i(\mathbf{s}) \leqslant 1 \ \forall \mathbf{s}. \end{aligned} \tag{7}$$

To get a better handle on the solution to the private value auction problem, we numerically computed solutions to (6) and (7) for examples with two agents. For the simulations, we considered analogues of these programs but with finitely many actions and signals.[27]

[27]With the lower bound on information, we have not proven (as we did without the lower bound) that dual reductions can be truncated with arbitrarily small loss. In the case of information structures, this is true for the same reason as it was before: the likelihood of finite length signal sequences must go to zero as the length goes to infinity. Hence, the change in the potential from censoring will be small as long as the censoring starts after a sufficiently large length. For mechanisms, it is more complicated, and we

For the first simulation, the prior is that all value profiles in $\{3, 4\}^2$ are equally likely. This case serves as a kind of sanity check, because we can easily argue that there is no duality gap and pin down the optimal guarantee and potential. Indeed, the potential is minimized by the independent private value information structure in which the agents learn their values and nothing else. In this case, the optimal auction would be to offer the good to one of the agents at a price of 4, and if they do not want it, offer it to the other agent at a price of 3. This generates a revenue of 3.5. And in fact, for every $\epsilon > 0$, there is a mechanism whose guarantee is at least $3.5 - \epsilon$: just perturb the aforementioned mechanism by offering the good at prices $4 - \epsilon$ and $3 - \epsilon$, so that reporting one's true value is the unique best response. Thus, in this case, there is no duality gap, and max guarantee is equal to min potential.

For our simulations for minimum potential, we allowed sequences up to length 8, and we set C to be 3. The simulated minimum upper bound on the potential for ordered information structures turns out to be *exactly* 3.5 (up to 8 decimal places). For simulations of the maximum guarantee, with sequences up to length 8 and C equal to 3, the simulated lower bound on the guarantee is approximately 3.269. When we use sequences up to length 150, with $C = 45$, the simulated lower bound increases to 3.483.[28] Thus, while we do not theoretically establish that the max guarantee with private values can be approached with finite ordered mechanisms, the loss in the guarantee from using such mechanisms is at most 0.48% percent of the optimal guarantee.

For our second simulation, the value profile is uniformly distributed on $\{(3, 3), (4, 3), (3, 4)\}$. Hence, values are correlated, and if agents only observed their own values, then it would be possible to extract all of the surplus, for an expected revenue of $11/3 \approx 3.667$. This cannot be the min potential (even making one of the agents' values public information would lead to a lower potential), but we do not have a theoretical argument that pins down the min potential or establishes whether or not there is a duality gap. However, a straightforward

do not have a theorem that the loss from restricting to finite dually reduced mechanisms can be made arbitrarily small. Without the lower bound, we were able to do so because, by the ergodic theorem, the mixture associated with the reduced action converged to a limit measure. But with the lower bound, it is not necessarily the case that the mixture converges as the length of the sequence goes to infinity, unless the distribution of signals in the sequence is itself ergodic. Nonetheless, the guarantees that we simulate with finite length sequences are a lower bound on the max guarantee.

[28]The number of binary sequences of length 150 is enormous; there is no hope of solving the associated linear program numerically, without a shortcut. In all of the simulations that we conducted with short sequences, of length less than 8, the optimal value was the same when we only allowed actions that were sequences of low values, or sequences of low values followed by a single high value. For the simulation with length 150 sequences, we only allowed this subset of action sequences. This still gives us a valid lower bound on the max guarantee, since it is equivalent to dropping equilibrium constraints associated with all other sequences in calculating a lower bound on the guarantee.

lower bound on the guarantee comes from the optimal dominant strategy mechanism[29], which in this case is the same as before: offer the good to one agent at a price of 4, and if they do not want it, sell it to the other agent at a price of 3, which yields revenue of $10/3 \approx 3.333$. As before, we could perturb this mechanism to make truthtelling a strictly dominant strategy, so that 10/3 is a lower bound on the max guarantee.

For the simulated min potential, we allowed sequences up to length 30 and set $C = 30$. The computed upper bound for the min potential is 3.479.[30] For the simulated max guarantee, we allowed sequences up to length 150 and set $C = 45$, and computed a lower bound for the max guarantee of approximately 3.371. Thus, a seller can achieve a strictly higher guarantee with a Bayesian mechanism than with a (perturbed) dominant strategy mechanism. This finding is consistent with the analysis of Chung and Ely (2007), who present examples in which no common prior information structure can rationalize a dominant strategy mechanism as being optimal. We have not pushed the limits of these simulations in terms of increasing the length of the sequences or optimizing C, and we have no reason to think that there is a duality gap.

The primary lesson from these simulations is a proof of concept, that the informationally robust approach can be operationalized when there are lower bounds on the agents' information. But there is clearly much more to do. What do the potential-minimizing information structures and guarantee-maximizing mechanisms look like? In the independent case, is the guarantee maximized by mechanisms that resemble perturbed dominant strategy mechanisms, or are they completely different? In the correlated case, what is the form of the extra information in the sequences, and what interim beliefs are induced that prohibit full surplus extraction? And how do the optimal ordered mechanisms leverage the correlation in private values in order to maximize the guarantee?

We leave the task of answering these tantalizing questions to the next installment in this research agenda. But it seems clear that these mechanisms represent a promising alternative to the "strange auctions" first described in Myerson (1981) for revenue generation with correlated private values. The theory also resolves the full-surplus extraction paradox, in a manner that is very much in the spirit of Bergemann and Morris (2005) and Chung and Ely (2007): If agents have correlated private values but can have aribtrary information, then

[29]Unlike Crémer and McLean (1988), we insist on the ex post individual rationality condition for a dominant strategy mechaism.

[30]A computational shortcut is that we restricted the signal sequences to be monotone. Again, this restriction would lead to an even more permissive upper bound on the min potential.

full surplus extraction may be impossible, and such additional information may rationalize more simple and portable mechanisms as being optimal.[31]

8 Conclusion

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. The key ideas also generalize to models where there are upper and lower bounds on the agents' information.

While the performance 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.

References

Bergemann, D., A. Bonatti, and A. Smolin (2018): "The design and price of information," *American economic review*, 108, 1–48.

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

[31]See Pusztai and Rahman (2025) for a related point that more information for the agents can prevent full surplus extraction.

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

——— (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.

——— (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.

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.

Crémer, J. and R. P. McLean (1988): "Full Extraction of the Surplus in Bayesian and Dominant Strategy Auctions," *Econometrica*, 1247–1257.

Lehrer, E., D. Rosenberg, and E. Shmaya (2013): "Garbling of signals and outcome equivalence," *Games and Economic Behavior*, 81, 179–191.

Loertscher, S. and E. V. Muir (2024): "Optimal Hotelling Auctions," Working paper.

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.

Pusztai, P. and D. M. Rahman (2025): "Redistributing Information Rents," *Available at SSRN 5160344*.

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

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

Yang, F. (2023): "Nested bundling," Tech. rep., Stanford University, working Paper.

A Omitted Proofs

Proof of Theorem 5. The first inequality follows from the fact that this is a Lagrangian relaxation of the guarantee program, where we have fixed a particular set of multipliers, as in Lemma 1. For the second inequality, we have that $\underline{G}_I(X, m^*, C)$ is equal to

$$\sum_{s,\theta}\sigma(s,\theta)\inf_{\mathbf{s}}\sum_{\omega}\left[w(\omega,\theta)m^*(\omega|\mathbf{s})+C\sum_{i,a}u_i(\omega,\theta)m(\omega|a)\left(b_i^*(a_i|\mathbf{s}_i,s_i)-b_i^*(a_i|\mathbf{s}_i)\right)b_{-i}^*(a_{-i}|\mathbf{s}_{-i})\right]$$

$$=\sum_{s,\theta}\sigma(s,\theta)\inf_{\mathbf{s}}\sum_{\omega}\Bigg[w(\omega,\theta)m^*(\omega|\mathbf{s})$$
$$+C\sum_{i,a}u_i(\omega,\theta)m(\omega|a)\left(\sum_{a_i'}\frac{\alpha^*(a_i|a_i',s_i)}{C}b_i^*(a_i'|\mathbf{s}_i)-b_i^*(a_i|\mathbf{s}_i)\right)b_{-i}^*(a_{-i}|\mathbf{s}_{-i})\Bigg]$$

$$=\sum_{s,\theta}\sigma(s,\theta)\inf_{\mathbf{s}}\sum_{\omega}\Bigg[w(\omega,\theta)m^*(\omega|\mathbf{s})$$
$$+C\sum_{i,a}u_i(\omega,\theta)\left(\sum_{a_i'}\frac{\alpha^*(a_i'|a_i,s_i)}{C}b_i^*(a_i|\mathbf{s}_i)m(\omega|a_i',a_{-i})-b_i^*(a_i|\mathbf{s}_i)m(\omega|a)\right)b_{-i}^*(a_{-i}|\mathbf{s}_{-i})\Bigg]$$

$$=\sum_{s,\theta}\sigma(s,\theta)\inf_{\mathbf{s}}\sum_{\omega}\left[w(\omega,\theta)m^*(\omega|\mathbf{s})+\sum_{i,a}u_i(\omega,\theta)\sum_{a_i'}\alpha^*(a_i'|a_i,s_i)\left(m(\omega|a_i',a_{-i})-m(\omega|a)\right)b^*(a|\mathbf{s})\right]$$

$$=\sum_{s,\theta}\sigma(s,\theta)\inf_{\mathbf{s}}\sum_{\omega,a}\left[w(\omega,\theta)m(\omega|a)+\sum_{i}u_i(\omega,\theta)\sum_{a_i'}\alpha^*(a_i'|a_i,s_i)\left(m(\omega|a_i',a_{-i})-m(\omega|a)\right)\right]b^*(a|\mathbf{s})$$

$$\geqslant \sum_{s,\theta} \sigma(s,\theta) \inf_a \sum_\omega \left[w(\omega,\theta) m(\omega|a) + \sum_i u_i(\omega,\theta) \sum_{a_i'} \alpha^*(a_i'|a_i,s_i) \left(m(\omega|a_i',a_{-i}) - m(\omega|a) \right) \right]$$
$$= G_I(M).$$

□

Proof of Theorem 6. As in the proof of Theorem 2, we will show that the lower bound $\overline{P}(\overline{X},\sigma^*,C)$ is the same as the optimal Lagrangian for I', but with an additional restriction on the set of mechanisms that are feasible:

$$\overline{P}(\overline{X},\sigma^*,C)$$
$$= \sum_{\mathbf{s}} \max_\omega \sum_{\theta,s'} \Bigg[w(\omega,\theta)\rho(\mathbf{s}|\mathbf{s}^0,s')\sigma'(\mathbf{s}^0,s',\theta)$$
$$+ C\sum_i u_i(\omega,\theta) \left(\rho_i(\mathbf{s}_i|\mathbf{s}_i^0,s_i')\sigma'(\mathbf{s}^0,s',\theta) - \sum_{s_i} \rho_i((s_i,\mathbf{s}_i)|s_i,s_i')\sigma'((s_i,\mathbf{s}_{-i}^0),s',\theta) \right) \rho_{-i}(\mathbf{s}_{-i}|\mathbf{s}_{-i}^0,s_{-i}') \Bigg]$$
$$= \sum_{\mathbf{s}} \max_\omega \sum_{\theta,s'} \Bigg[w(\omega,\theta)\rho(\mathbf{s}|\mathbf{s}^0,s')\sigma'(\mathbf{s}^0,s',\theta)$$
$$+ C\sum_i u_i(\omega,\theta) \Bigg(\rho_i(\mathbf{s}_i|\mathbf{s}_i^0,s_i')\sigma'(\mathbf{s}^0,s',\theta)$$
$$- \sum_{s_i,\widehat{s}_i'} \rho_i(\mathbf{s}_i|\mathbf{s}_i^0,\widehat{s}_i') \frac{\alpha_i^*(\mathbf{s}_i^0,\widehat{s}_i'|s_i,s_i')}{C} \sigma'((s_i,\mathbf{s}_{-i}^0),s',\theta) \Bigg) \rho_{-i}(\mathbf{s}_{-i}|\mathbf{s}_{-i}^0,s_{-i}') \Bigg]$$
$$= \sum_{\mathbf{s}} \max_\omega \sum_{\theta,s'} \Bigg[w(\omega,\theta)\rho(\mathbf{s}|\mathbf{s}^0,s')\sigma'(\mathbf{s}^0,s',\theta)$$
$$+ C\sum_i u_i(\omega,\theta) \Bigg(\rho_i(\mathbf{s}_i|\mathbf{s}_i^0,s_i')\sigma'(\mathbf{s}^0,s',\theta)$$
$$- \sum_{s_i,\widehat{s}_i'} \rho_i(\mathbf{s}_i|\mathbf{s}_i^0,s_i') \frac{\alpha_i^*(\mathbf{s}_i^0,s_i'|s_i,\widehat{s}_i')}{C} \sigma'((s_i,\mathbf{s}_{-i}^0),\widehat{s}_i',s_{-i}',\theta) \Bigg) \rho_{-i}(\mathbf{s}_{-i}|\mathbf{s}_{-i}^0,s_{-i}') \Bigg]$$
$$= \sum_{\mathbf{s}} \max_\omega \sum_{\theta,s'} \Bigg[w(\omega,\theta)\sigma'(\mathbf{s}^0,s',\theta)$$
$$+ \sum_i u_i(\omega,\theta) \left(C\sigma'(\mathbf{s}^0,s',\theta) - \sum_{s_i,\widehat{s}_i'} \alpha_i^*(\mathbf{s}_i^0,s_i'|s_i,\widehat{s}_i')\sigma'((s_i,\mathbf{s}_{-i}^0),\widehat{s}_i',s_{-i}',\theta) \right) \Bigg] \rho(\mathbf{s}|\mathbf{s}^0,s')$$
$$= \sum_{\mathbf{s}} \max_\omega \sum_{\theta,s'} \Bigg[w(\omega,\theta)\sigma'(\mathbf{s}^0,s',\theta) + \sum_i \beta_i^*(\mathbf{s}_i^0,s_i')u_i(\omega,\theta)\sigma'(\mathbf{s}^0,s',\theta)$$

$$+ \sum_{i,s_i,\widehat{s}'_i} u_i(\omega,\theta)\left(\alpha_i^*(s_i,\widehat{s}'_i|\mathbf{s}_i^0,s'_i)\sigma'(\mathbf{s}^0,s',\theta) - \alpha_i^*(\mathbf{s}_i^0,s'_i|s_i,\widehat{s}'_i)\sigma'((s_i,\mathbf{s}_{-i}^0),\widehat{s}'_i,s'_{-i},\theta)\right)\Bigg]\rho(\mathbf{s}|\mathbf{s}^0,s'),$$

which is the original Lagrangian, as desired, except that now the mechanism is mediated through ρ. It is clearly smaller than

$$\sum_{\mathbf{s},s'}\max_{\omega}\sum_{\theta}\Bigg[w(\omega,\theta)\sigma'(\mathbf{s}^0,s',\theta) + \sum_i \beta_i^*(\mathbf{s}_i^0,s'_i)u_i(\omega,\theta)\sigma'(\mathbf{s}^0,s',\theta)$$
$$+ \sum_{i,s_i,\widehat{s}'_i} u_i(\omega,\theta)\left(\alpha_i^*(s_i,\widehat{s}'_i|\mathbf{s}_i^0,s'_i)\sigma'(\mathbf{s}^0,s',\theta) - \alpha_i^*(\mathbf{s}_i^0,s'_i|s_i,\widehat{s}'_i)\sigma'((s_i,\mathbf{s}_{-i}^0),\widehat{s}'_i,s'_{-i},\theta)\right)\Bigg]\rho(\mathbf{s}|\mathbf{s}^0,s')$$
$$= \sum_{\mathbf{s}^0,s'}\max_{\omega}\sum_{\theta}\Bigg[w(\omega,\theta)\sigma'(\mathbf{s}^0,s',\theta) + \sum_i \beta_i^*(\mathbf{s}_i^0,s'_i)u_i(\omega,\theta)\sigma'(\mathbf{s}^0,s',\theta)$$
$$+ \sum_{i,s_i,\widehat{s}'_i} u_i(\omega,\theta)\left(\alpha_i^*(s_i,\widehat{s}'_i|\mathbf{s}_i^0,s'_i)\sigma'(\mathbf{s}^0,s',\theta) - \alpha_i^*(\mathbf{s}_i^0,s'_i|s_i,\widehat{s}'_i)\sigma'((s_i,\mathbf{s}_{-i}^0),\widehat{s}'_i,s'_{-i},\theta)\right)\Bigg]$$
$$= P(I').$$

□

B Dual Reductions for Optimal Auctions

For the optimal auctions problem, the dual reduction I^* of a finite information structure I is defined as in Section 5, using the optimal Lagrange multipliers α^* and β^* for the revenue maximization program.

Theorem 7. *For the optimal auctions problem, given any finite information structure $I = (S,\sigma)$ and corresponding dual reduction $(\overline{X},\sigma^*)$ and multiplier C, we have that $P(\overline{X},\sigma^*) \leqslant \overline{P}(\overline{X},\sigma^*,C) \leqslant P(I)$.*

Proof. Applying Theorem 3 in Appendix B.2 of Brooks and Du (2024), we have that the potential is at most the first-order upper bound $\overline{P}(I^*,C)$. This is

$$\overline{P}(\overline{X},\sigma^*)$$
$$= \sum_x \max_{q,t}\sum_\theta\Bigg[(\sum_i t_i)\sigma^*(x,\theta) - C\sum_i(q_i\theta_i - t_i)\left(\sigma^*(x_i+1,x_{-i},\theta) - \sigma^*(x,\theta)\right)\Bigg]$$
$$= \sum_x \max_{q,t}\sum_\theta\Bigg[(\sum_i t_i)\sigma^*(x,\theta) - C\sum_s \sigma(s,\theta)\sum_i(q_i\theta_i - t_i)\left(\rho_i(x_i+1|s_i) - \rho_i(x_i|s_i)\right)\rho_{-i}(x_{-i}|s_{-i})\Bigg]$$

$$
\begin{aligned}
&= \sum_x \max_{q,t} \sum_{s,\theta} \Bigg[(\sum_i t_i)\sigma(s,\theta)\rho(x|s) \\
&\qquad - C\sigma(s,\theta)\sum_i (q_i\theta_i - t_i)\left(\sum_{s_i'} \frac{\alpha_i^*(s_i'|s_i)\rho(x_i|s_i')}{C} - \rho_i(x_i|s_i)\right)\rho_{-i}(x_{-i}|s_{-i})\Bigg] \\
&= \sum_x \max_{q,t} \sum_{s,\theta} \Bigg[(\sum_i t_i)\sigma(s,\theta)\rho(x|s) \\
&\qquad - \sigma(s,\theta)\sum_i (q_i\theta_i - t_i)\Bigg(\sum_{s_i'} \alpha_i^*(s_i'|s_i)\rho(x_i|s_i') \\
&\qquad\qquad - \left(\sum_{s_i'} \alpha_i^*(s_i'|s_i) + \beta_i^*(s_i)\right)\rho_i(x_i|s_i)\Bigg)\rho_{-i}(x_{-i}|s_i)\Bigg] \\
&= \sum_x \max_{q,t} \sum_{s,\theta} \Bigg[(\sum_i t_i)\sigma(s,\theta)\rho(x|s) + \sum_i (q_i\theta_i - t_i)\beta_i^*(s_i)\sigma(s,\theta)\rho(x|s) \\
&\qquad - \sum_i (q_i\theta_i - t_i)\left(\sum_{s_i'} \alpha_i^*(s_i|s_i')\sigma(s_i', s_{-i}, \theta) - \sum_{s_i'} \alpha_i^*(s_i'|s_i)\sigma(s,\theta)\right)\rho(x|s)\Bigg].
\end{aligned}
$$

In each maximization over (q,t), we allow $t \in \mathbb{R}^N$ and $q \in \Delta(\{0,1,\ldots,N\})$.

The above equation is clearly weakly less than

$$
\begin{aligned}
&\sum_{x,s} \max_{q,t} \sum_{\theta} \Bigg[(\sum_i t_i)\sigma(s,\theta)\rho(x|s) + \sum_i (q_i\theta_i - t_i)\beta_i^*(s_i)\sigma(s,\theta)\rho(x|s) \\
&\qquad - \sum_i (q_i\theta_i - t_i)\left(\sum_{s_i'} \alpha_i^*(s_i|s_i')\sigma(s_i', s_{-i}, \theta) - \sum_{s_i'} \alpha_i^*(s_i'|s_i)\sigma(s,\theta)\right)\rho(x|s)\Bigg] \\
&= \sum_s \max_{q,t} \sum_{\theta,x} \Bigg[(\sum_i t_i)\sigma(s,\theta)\rho(x|s) + \sum_i u_i(\omega,\theta)\beta_i^*(s_i)\sigma(s,\theta)\rho(x|s) \\
&\qquad - \sum_i (q_i\theta_i - t_i)\left(\sum_{s_i'} \alpha_i^*(s_i|s_i')\sigma(s_i', s_{-i}, \theta) - \sum_{s_i'} \alpha_i^*(s_i'|s_i)\sigma(s,\theta)\right)\rho(x|s)\Bigg] \\
&= P(S,\sigma).
\end{aligned}
$$

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^*.

□