



EFFICIENT RANDOM ASSIGNMENT WITH CARDINAL AND ORDINAL PREFERENCES

James Fisher

United States Automobile Association (USAA), USA

jamescdfisher@gmail.com

ABSTRACT

We develop a finite random assignment model where players know either their cardinal or their ordinal preferences and may make cardinal or ordinal reports to an assignment mechanism. Under truthful reporting, we find that all mechanisms that disregard the cardinal information in players' reports (e.g., Deferred Acceptance and Probabilistic Serial) are utilitarian inefficient, as are classic mechanisms that make use of cardinal information (e.g., Pseudo-markets). Motivated by these negative results, we introduce a "Simple Mechanism" that makes use of cardinal information "in the right way." We establish that this mechanism is utilitarian efficient, treats equals equally, and makes truth-telling almost Bayesian incentive compatible.

Keywords: Cardinal and ordinal preferences, random assignment, utilitarian efficiency.

JEL Classification Numbers: C7, D7, D8.

1. INTRODUCTION

Consider the problem of assigning two equally qualified doctors, Ann and Bob, to one of two medical research fellowships, at UC San Francisco and Mount Sinai Hospital in New York. Ann has prior knowledge of the colleagues at

I am grateful to Isa Hafalir, Tim Flannery, Antonio Miralles, Marek Pycia, John Wooders, Tiemen Woutersen, Yilei Zhang, and several anonymous referees for helpful comments, discussions, and insights. This paper was partially written while at University of California Los Angeles and the University of Arizona; I am grateful for their hospitality. All remaining errors are my own.

each hospital, the location (and available activities), and the salary, and has concluded that her value for UC San Francisco is 2, while her value for Mount Sinai is 1. Bob, who has been traveling with Doctors without Borders for the last few years, only knows that he strictly prefers UC San Francisco to Mount Sinai. The question is: given these preferences, how should we assign the fellowships to maximize social surplus/utilitarian welfare (without using money)?¹

Our aim is to answer questions of this type. The key insight of our approach is that Bob, like every other person, has a preference intensity, i.e., he can like UC San Francisco a bit more than Mount Sinai or he can like it a lot more. Thus, Bob's ordinal preference is actually induced by a true, underlying cardinal preference that, we suppose, he does not know. (We discuss the reasons why a person may only know their ordinal preferences below.)

Since Bob has an underlying cardinal preference, we are able to estimate it from his ordinal preference and are able to use this estimate to compute a surplus maximizing assignment. For instance, if he has three possible values for each hospital of 1 , $\frac{3}{2}$, and 2 , which are equally likely, then we estimate that his values for UC San Francisco and Mount Sinai are $\frac{11}{6}$ and $\frac{7}{6}$ respectively.² It follows that it is surplus maximizing to send Ann to San Francisco and send Bob to New York – this assignment generates an (expected) social surplus of $2 + \frac{7}{6} = \frac{19}{6}$.

In this paper, we study the elicitation and estimation of preferences, as well as their use in constructing surplus maximizing assignments, in a simple random assignment model with a finite number of players N and a finite number of objects K . Each player may get at most one object and an assignment is a map that tells us which players get which objects; a mixed assignment is a lottery over assignments.

The players are Bayesians with Von-Neumann Morgenstern utilities. To capture the idea that some know their cardinal preferences and others only

¹ This paper contributes to the “random assignment literature,” which seeks to understand how goods should be allocated when monetary transfers are not possible. This is the case, for instance, with dorm rooms, fellowships, organs, school seats, and public programs (e.g., day-care and low-cost housing) where resell is impossible and other transfers to prevent participation in the assignment and/or alter reports are illegal. (In fact, California's Cartwright Act would prohibit Bob from paying Ann to not apply to UC San Francisco.)

² There are nine equally likely possible pairs of values for UC San Francisco and Mount Sinai: $(1, 1)$, $(1, \frac{3}{2})$, \dots , $(2, \frac{3}{2})$, and $(2, 2)$. Of these, only $(\frac{3}{2}, 1)$, $(2, 1)$, and $(2, \frac{3}{2})$ are compatible with Bob's ordinal preference. Averaging these three gives $(\frac{11}{6}, \frac{7}{6})$.

know their ordinal preferences, we suppose that nature draws a vector of “true values” for each player from a finite set of possible true value vectors. Nature then randomly determines which players observe their true values and which observe the rankings implied by their true values.

To digress, there are a number of reasons a person may only observe her ordinal preferences. For example, it may be that the objects are experience goods and, absent prior consumption, she may only gain a general ranking from the reviews and recommendations of others: a classic example is restaurants, where a Zagat’s guide or Yelp reviews only provide a general ranking and the actual dining experiences are what provide the true values; additional examples include colleges, housing, and the like. Alternatively, the objects may be complex and hard to evaluate – e.g., home mortgages, health insurance contracts, employment opportunities, retirement plans, and the like. Hence, a person may avail herself of the reviews and recommendations of others to gain a general idea of how the objects rank or, if she has access to the needed information and is sufficiently skilled, she may be able to determine her cardinal values.³ We abstract away from specifics by assuming some players are endowed with knowledge of their cardinal preferences, while others are not.

After players learn their true values or rankings, i.e., their types, they send reports to a mechanism. The mechanism takes in these reports, processes them, and then returns a mixed assignment. Subsequently, an outcome of this assignment is realized and players receive their objects. We focus on direct mechanisms for simplicity.

We are primarily concerned with interim (utilitarian) efficiency. An assignment is interim efficient if it maximizes (expected) social surplus once players know their types. A mechanism is interim efficient if it selects an interim efficient assignment (when players report truthfully).⁴ Interim efficiency implies interim and ex-post Pareto efficiency, as well as interim and ex-post individual rationality.

Within this framework, we first examine whether classical mechanisms –

³ Introspection suggests that, as a result of general reviews and rankings, a person may only get an “imprecise” estimate of her preference intensity – e.g., she may only know whether one object is a bit more valuable than another or a lot more valuable – which she is unable to refine to a point estimate. In contrast, a Bayesian player always has a point estimate of her expected preference intensity. Hence, we view the Bayesian nature of players as an analytically tractable approximation of real-world estimating behavior.

⁴ We study incentive compatibility separately from efficiency.

e.g., Boston, Deferred Acceptance, Random Serial Dictatorship, Probabilistic Serial, and Pseudo-markets – produce efficient assignments. Some of these mechanisms make use of the cardinal information contained in players’ reports (e.g., Pseudo-markets), while others “disregard” it (e.g., Boston, Deferred Acceptance, Random Serial Dictatorship, and Probabilistic Serial).

Our first principal result is a negative one: under broad conditions – roughly, that two different, possible vectors of true values induce the same ordinal preference – any mechanism that disregards cardinal information (i.e., that only uses the ranking implied by players’ reports) is interim inefficient (see Proposition 1). This result strengthens the classical intuition that such mechanisms “may be inefficient” and implies that they only achieve the second best.⁵ The result is grounded in the fact that, when a mechanism disregards cardinal information, objects do not go to those who value them most. (We develop this intuition and those related to our other results in detail in the Illustration of Section 3.) Subsequently, we investigate whether mechanisms that make use of cardinal information are interim efficient and find that inefficiencies occur (see Examples 1 and 2).

Motivated by these negative findings, we introduce a “Simple Mechanism.” It takes in players’ ordinal and cardinal reports and then constructs estimates of their true values by treating their reports as though they were their types. It next computes σ , the set of pure assignments that maximize the sum of players’ estimated payoffs. Subsequently, it generates a mixed assignment by implementing each element of σ with equal probability.

Our second principal result establishes that the Simple Mechanism is interim efficient and “symmetric” (see Proposition 2). The intuition for efficiency is that, under truthful reporting, the mechanism correctly estimates (expected) true values and thus maximizes social surplus. For instance, if we run it in our fellowship problem, then it generates the efficient assignment – sending Ann to San Francisco and Bob to New York – because it estimates that she values San Francisco more than he does. The Simple Mechanism’s symmetry ensures that “equals are treated equally,” i.e., if two players make the same report and have the same type, then they have the same payoff. Symmetry follows from

⁵ The magnitude of the inefficiencies can be substantial. For instance, if we run Random Serial Dictatorship for our fellowship problem, then Ann and Bob would both have a $\frac{1}{2}$ chance of ending up at either fellowship; social surplus in this assignment is only $(2+1)\frac{1}{2} + (\frac{11}{6} + \frac{7}{6})\frac{1}{2} = \frac{18}{6}$. Thus, the inefficiency is $\frac{1}{6}$, which is approximately 25% of the Ann and Bob’s average value for the fellowships.

the mechanism's uniform randomization.

Since efficiency implies individual rationality, it follows that players are open to using the Simple Mechanism. The real question, then, is whether they report truthfully so that the mechanism achieves its designed ends.

Our third principal result tackles this question: we establish that the Simple Mechanism makes truth-telling (i.e., truthful reporting of one's type) "approximately" incentive compatible. Specifically, we prove that for every $\varepsilon > 0$, no player i gains more than ε by lying strategically about her type to the Simple Mechanism instead of truthfully reporting it when (i) all other players report truthfully and (ii) and the number of players N is large (see Proposition 3). The core intuition is that, holding the set of objects fixed, as N grows, the probability i can get any object decreases because of "competition" with the other players. Thus, i 's payoffs to lying and to telling the truth converge to her value of being unassigned from above, implying her gain from lying is less than ε for large N .

While there are situations where supply of objects is fixed, it is often natural to think that it increases with the number of players – e.g., colleges offer more courses in response to rising enrollments. Our fourth principal result addresses this point: we establish that there is a way to grow the set of objects with N via replication, while maintaining approximate incentive compatibility (see Proposition 4). The key insight is that, as N grows, replication must happen sufficiently slowly so that competitive pressures are preserved.

Since strategic manipulation is often costly – in terms of effort and time – our results imply that the Simple Mechanism makes truth-telling a Nash equilibrium. This, along with efficiency and symmetry, suggest that the Simple Mechanism may be a good mechanism to employ in situations where there is substantial demand for objects – as might be the case in the assignment of prime university housing or courses, in the awarding of grants or fellowships, in the allocation of slots in programs for free/subsidized day-care or housing, in the apportionment of prime employee parking spots or offices, or even in the distribution of kidneys and other organs.

The rest of this section surveys the related literature. Section 2 describes the model. Section 3 illustrates our results via an example and discusses the intuition behind them. Section 4 gives Proposition 1 and Examples 1 and 2. Section 5 describes the Simple Mechanism and gives Propositions 2 to 4. All proofs are provided in the Appendix and the Online Appendix collects supplemental results.

RELATED LITERATURE

We make three contributions to the modern random assignment literature. First, we allow players to know either their ordinal or cardinal preferences. Second, we strengthen a classic intuition by showing that all mechanisms that disregard cardinal information must be interim inefficient. Third, and most importantly, we show that efficiency can be recovered in an approximately incentive compatible fashion by giving a new, Simple Mechanism.

Regarding the first contribution, in the random assignment literature either (i) all players know their own cardinal preferences (e.g., [Bu \(2014\)](#), [Budish \(2011\)](#), [Hafalir & Miralles \(2014\)](#), [He et al. \(2018\)](#), and [Hylland & Zeckhauser \(1979\)](#)) or (ii) all players know their own ordinal preferences (e.g., [Abdulkadiroglu & Sonmez \(1998\)](#), [Bogomolnaia & Moulin \(2001\)](#), [Kasajima \(2013\)](#), [Liu & Pycia \(2016\)](#), [Nicolo & Rodriguez-Alvarez \(2013\)](#), and [Sonmez & Unver \(2010\)](#)). We view this polarity as artificial – since, in practice, some people know more about their preferences than others – and we show one way to dispense with it.

Regarding the second contribution, there is a classic intuition in the random assignment literature: mechanisms that disregard cardinal information *may* be (interim) inefficient (e.g., [Abdulkadiroglu et al. \(2008\)](#)). We strengthen this intuition by giving broad conditions under which these mechanisms *must be* inefficient. In this capacity, our work is closely related to the work of [Pycia \(2011\)](#). Pycia establishes that any ordinal, symmetric, regular, Pareto efficient, and asymptotically strategy-proof mechanism is inefficient.⁶ Our relative contribution is to remove many of these qualifiers and cleanly identify the source of the inefficiency.

Our (interim) inefficiency result complement those of [Che & Tercieux \(2013\)](#) and [Zhou \(1990\)](#). Che and Tercieux show that the ex-ante per-capita utilitarian inefficiency of Pareto efficient mechanisms converges to zero as the number of players (and objects) goes to infinity. In light of this, our inefficiency result implies that the ex-ante per-capita utilitarian inefficiency must be strictly positive outside of the limit when a mechanism disregards

⁶ An ordinal mechanism is one that only accepts ordinal reports, while a regular mechanism is one where, when the number of players is large, no small group of players can “substantially impact the allocations of other...” players ([Liu & Pycia, 2016](#), 3). In addition, a mechanism is asymptotically strategy-proof when, for each $\epsilon > 0$, players do not gain for than ϵ from lying when (i) all other players tell the truth, (ii) all players’ types are known, and (iii) the number of players is large.

cardinal information. Zhou identifies a parallel source of inefficiency: the interaction of strategy-proofness and symmetry. Together, our results enrich the literature’s understanding of when inefficiencies occur.

Regarding the third contribution, the Simple Mechanism is most closely related to the work of [Nguyen et al. \(2014\)](#). Nguyen et al. develop a random assignment model where players can acquire more than one object. They assume that players know their cardinal preferences and develop a mechanism based on a linear programming problem that maximizes social surplus subject to the constraint that the assignment be envy-free. They show that their mechanism is (approximately) constrained efficient and is asymptotically strategy-proof when the programming problem has a *unique* solution. Our relative contribution is to give the Simple Mechanism, which is always unconstrained efficient and always has economically compelling truth-telling properties.⁷

The Simple Mechanism is further related to the work of [Budish \(2011\)](#). Budish studies (in passing) a “utilitarian” mechanism that solicits players cardinal reports and assigns objects to maximize utilitarian surplus, using an environment where cardinal preferences are known.⁸ Unlike the Simple mechanism, Budish’s mechanism does not randomize. He concludes that this mechanism is “manipulable” in the sense that a player may always have an incentive to lie; this conclusion extends to the Simple Mechanism (and we discuss it in the Illustration of Section 3). Our relative contribution is to show that the gains to manipulation shrink as competition intensifies due to growth in the number of players.

2. THE MODEL

To rigorously study efficient assignment, we need to pin down the economic environment and introduce a few definitions concerning efficiency and incentive compatibility.

ECONOMIC ENVIRONMENT

There is a finite set of players $\mathcal{N} = \{1, \dots, N\}$ and a finite set of objects $\mathcal{O} = \{1, \dots, K\}$, where N and K are positive integers. We write i for the i -th player and o for the o -th object. Each player may be assigned to an object or

⁷ Note, however, that the Simple Mechanism may not be envy-free – e.g., in the efficient assignment of our fellowship example, Bob wants Ann’s position.

⁸ The thrust of [Budish \(2011\)](#) is the introduction of an “approximate Pseudo-market” mechanism, which is Pareto efficient and “almost envy-free,” but need not be interim efficient or symmetric.

may be unassigned, and each object may be assigned to at most one player. We adopt the convention that a player is unassigned if she is assigned to the “null” item η and write $\mathcal{O}' = \mathcal{O} \cup \{\eta\}$. A **pure assignment** ϕ is a function that specifies each player’s assignment, i.e., $\phi : \mathcal{N} \rightarrow \mathcal{O}'$ such that, for each $i \in \mathcal{N}$ and each $o \in \mathcal{O}$, $\phi(i) = o \iff \phi^{-1}(o) = i$, where ϕ^{-1} is the inverse of ϕ .

We write Φ for the finite set of all pure assignments and we index the elements of this set from 1 to $|\Phi|$, where $|\cdot|$ gives the cardinality of a set. Let $\Delta\Phi = \{(x_1, \dots, x_{|\Phi|}) \in \mathbb{R}_+^{|\Phi|} \mid \sum_{j=1}^{|\Phi|} x_j = 1\}$ be the set of probability distributions on Φ . A $\tilde{\phi} = (\tilde{\phi}_1, \dots, \tilde{\phi}_{|\Phi|}) \in \Delta\Phi$ specifies a lottery over pure assignments where the l -th pure assignment is realized with probability $\tilde{\phi}_l$; the lottery is independent of all other randomness. We refer to $\tilde{\phi}$ as a **mixed assignment**. The **support** of $\tilde{\phi} = (\tilde{\phi}_1, \dots, \tilde{\phi}_j, \dots, \tilde{\phi}_{|\Phi|})$ is $\{\phi_j \in \Phi \mid \tilde{\phi}_j > 0\}$, i.e., the set of pure assignments on which it places non-zero probability.

We have in mind that some players know their cardinal values for objects and for being unassigned, while other players only know their ordinal preferences. To model this, we suppose that each player has a “true (monetary) value” for each object and for being unassigned. We only allow certain players to observe their true values. The other players observe binary relations that are “consistent” with their true values. Nature randomly decides which players observe their true values. For simplicity, the true values are independently and identically distributed and each player has the same chance of observing her true values.⁹

Specifically, for each player i , let $v_i = (v_{\eta i}, v_{1i}, v_{2i}, \dots, v_{Ki})$ denote the vector of i ’s true values over η and the objects in \mathcal{O} , where $v_{\eta i}$ is i ’s value of η , v_{1i} is i ’s value of object 1, v_{2i} is i ’s value of object 2, and so on. These values vectors are drawn by nature from a finite set $\mathcal{V} \subset \mathbb{R}^{K+1}$ according to probability mass function $f_v(v_i)$, where $f_v : \mathcal{V} \rightarrow (0, 1]$ such that $\sum_{v \in \mathcal{V}} f_v(v) = 1$; the draw is independent of all other randomness. These values are monetary in nature and are nontransferable (due to institutional, legal, ethical, or other constraints).

Let \mathcal{B} be the set of all complete, transitive, and reflexive binary relations on \mathcal{O}' . We say that a $\preceq \in \mathcal{B}$ is **consistent** with a $v = (v_\eta, v_1, v_2, \dots, v_K)$ if $v_o \leq v_{o'} \iff o \preceq o'$ and $v_o \not\leq v_{o'} \iff o \not\preceq o'$ for all o and o' in \mathcal{O}' . That is,

⁹ In the Online Appendix, we show that our results generalize to an environment where players have correlated true values, as well as arbitrary (and possibly correlated) knowledge of their true values.

\preceq is consistent with v if it is the ordinal preference of a decision maker whose true values are v . We also say that v is consistent with \preceq if \preceq is consistent with v . Let $\mathcal{B}' = \{\preceq \in \mathcal{B} \mid \preceq \text{ is consistent with a } v \in \mathcal{V}\}$ be the subset of \mathcal{B} whose binary relations are consistent with the values in \mathcal{V} ; we work principally with \mathcal{B}' .

Lemma 1. The Consistency Relationship.

For each $v \in \mathcal{V}$, there is a unique consistent $\preceq \in \mathcal{B}'(v)$. However, for each $\preceq \in \mathcal{B}'$, there may be multiple consistent valuations in \mathcal{V} .

For each player i , let \preceq_i denote the element of $\mathcal{B}'(v)$ that is consistent with her true values v_i . After nature draws v_i , nature either reveals it to player i with probability $\alpha \in [0, 1]$ or nature reveals \preceq_i with probability $1 - \alpha$; the lottery is independent of all other randomness. If i observes v_i , then she knows her cardinal preferences; however, if she observes \preceq_i , then she only knows her ordinal preferences.

Let θ_i denote nature's report to player i , which we refer to i 's **type**. Let $\Theta = \mathcal{V} \cup \mathcal{B}'$ denote the set of types. The probability i gets type θ_i is,¹⁰

$$\Pr(\theta_i) = \begin{cases} \alpha f_v(\theta_i) & \text{if } \theta_i \in \mathcal{V} \\ (1 - \alpha) \sum_{v \in I(\theta_i)} f_v(v) & \text{if } \theta_i \in \mathcal{B}', \end{cases} \quad (1)$$

where $I(\theta_i) = \{v \in \mathcal{V} \mid v \text{ is consistent with } \theta_i\}$. Let $\theta = (\theta_1, \dots, \theta_N)$ denote the profile of players' types. Due to independence, the probability of θ is $\Pr(\theta) = \prod_{i=1}^N \Pr(\theta_i)$. When $\alpha \in (0, 1)$, all type profiles occur with strictly positive probability because $f_v > 0$.

After each player observes her type, she makes a report about her preferences to an assignment mechanism, which processes the reports and generates a (mixed) assignment; we focus on direct mechanisms for simplicity. Let $r_i \in \Theta$ denote the report of player i and $r = (r_1, \dots, r_N)$ denote the profile of players' reports. Formally, a **mechanism** is a $M : \Theta^N \rightarrow \Delta\Phi$. After the mechanism generates its mixed assignment, the outcome of the assignment is realized and players receive their objects and their corresponding true values.

The players are Bayesians who have Von-Neumann Morgenstern utility. Thus, when player i gets type θ_i and makes her report, her expected true value

¹⁰ If $\theta_i \in \mathcal{V}$, then nature must have drawn θ_i for player i and revealed it to her, an event of probability $\alpha f_v(\theta_i)$. If $\theta_i \in \mathcal{B}'$, then nature must have drawn a vector v of true values that are consistent with θ_i and must not have revealed v to i , an event of probability $(1 - \alpha) \sum_{v \in I(\theta_i)} f_v(v)$.

vector is,¹¹

$$v_i^\dagger = (v_{\eta i}^\dagger, v_{1i}^\dagger, \dots, v_{oi}^\dagger, \dots, v_{Ki}^\dagger) = \begin{cases} \theta_i & \text{if } \theta_i \in \mathcal{V} \\ \sum_{v \in I(\theta_i)} v \frac{f_v(v)}{\sum_{v' \in I(\theta_i)} f_{v'}(v')} & \text{if } \theta_i \in \mathcal{B}'. \end{cases} \quad (2)$$

And her payoff, denoted $u_i(\cdot | \theta_i)$, is,

$$u_i(o | \theta_i) = v_{oi}^\dagger \text{ for each } o \in \mathcal{O}'.$$

In abuses of notation, (i) we write $u_i(\phi | \theta_i) = u_i(\phi(i) | \theta_i)$ for i 's payoff to the pure assignment ϕ and (ii) we write $u_i(\tilde{\phi} | \theta_i) = \sum_{j=1}^{|\Phi|} u_i(\phi_j | \theta_i) \tilde{\phi}_j$ for i 's payoff to in the mixed assignment $\tilde{\phi} = (\tilde{\phi}_1, \dots, \tilde{\phi}_{|\Phi|})$.

EFFICIENCY AND TRUTHFUL REPORTING

We introduce three definitions to formalize our efficiency and incentive compatibility notions. Our first pair of definitions concerns the efficiency of assignments and mechanisms.

Given a profile of players' types $\theta = (\theta_1, \dots, \theta_N) \in \Theta^N$, we say that the (mixed) assignment $\tilde{\phi}^*$ is **interim efficient** if it solves $\max_{\tilde{\phi} \in \Delta \Phi} \sum_{i \in \mathcal{N}} u_i(\tilde{\phi} | \theta_i)$, i.e., if it maximizes utilitarian welfare after nature has endowed types but before all uncertainty about true values has been resolved. If $\tilde{\phi}^*$ is not interim efficient, we say it is **interim inefficient**.

An interim efficient (mixed) assignment $\tilde{\phi}^*$ is also “interim” (“ex-post”) Pareto efficient, i.e., there is no other pure assignment that leaves all players weakly better off and one player strictly better off before (after) all uncertainty about true values has been resolved.¹² Interim and ex-post Pareto efficiency imply that $\tilde{\phi}^*$ is also “individually rational,” i.e., each player gets an element of \mathcal{O}' she likes at least as much as η both before and after all uncertainty about true values is resolved.¹³ In addition, when \mathcal{V} contains only “strict value

¹¹ There are only two types of zero probability events, (i) $\alpha = 0$ and i observes a type $\theta_i \in \mathcal{V}$, and (ii) $\alpha = 1$ and i observes a type $\theta_i \in \mathcal{B}'$. If (i), then we assume that θ is actually i 's vector of true values. If (ii), then we assume that i 's true value is drawn from $I(\theta)$ with probability $\frac{f_v(v)}{\sum_{v' \in I(\theta)} f_{v'}(v')}$ for each v therein. The display equation embeds both of these natural assumptions.

¹² It is evident that interim efficient implies interim Pareto efficiency; we establish that it also implies ex-post Pareto efficiency in the Online Appendix.

¹³ The individual rationality criterion reflects the idea that each player does weakly better by participating in the assignment instead of opting out and getting η with certainty.

vectors,¹⁴ then $\tilde{\phi}^*$ is “ordinally efficient” by [Bogomolnaia & Moulin \(2001\)](#) Lemma 2.

Since payoffs are continuous in the mixed assignment and the set of mixed assignments is compact, the Extreme Value Theorem applies and guarantees the existence of an interim efficient assignment. We have thus proved the following result.

Lemma 2. Existence of an Interim Efficient Assignment.

For each $\theta \in \Theta^N$, there is an interim efficient mixed assignment $\tilde{\phi}^$.*

We say that a mechanism M is **interim efficient** if $M(\theta)$ is an interim efficient (mixed) assignment for each profile of types $\theta \in \Theta^N$. In other words, an interim efficient mechanism is one that always selects an interim efficient assignment when the players truthfully report their types. Such a mechanism inherits the interim and ex-post Pareto efficiency of interim efficient assignments, as well as their ordinal efficiency and individual rationality. If M not interim efficient, we say that M is **interim inefficient**.

Our third definition concerns players’ incentives to truthfully report their types to a mechanism. For player i , let $U_i^M(r_i|\theta_i)$ denote i ’s payoff to reporting $r_i \in \Theta$ to the mechanism M given her type θ_i , when all other players truthfully report their types $\theta_{-i} = (\theta_1, \dots, \theta_{i-1}, \theta_{i+1}, \dots, \theta_N)$, i.e.,

$$U_i^M(r_i|\theta_i) = \sum_{\theta_{-i} \in \Theta^{N-1}} u_i(M(r_i, \theta_{-i})|\theta_i) \Pr(\theta_{-i}),$$

where $\Pr(\theta_{-i}) = \prod_{j \neq i} \Pr(\theta_j)$ is the probability of θ_{-i} .

For each $\varepsilon > 0$, a mechanism M is **ε -Bayesian incentive compatible** if, for each player i and each type $\theta_i \in \Theta$, i does not gain more than ε from lying strategically when everyone else tells the truth, i.e., if $\max_{r_i \in \Theta} U_i^M(r_i|\theta_i) \leq U_i^M(\theta_i|\theta_i) + \varepsilon$.¹⁵ Since it is often costly – in terms of effort and time – for a person to formulate a strategic manipulation, ε -Bayesian incentive compatibility is a compelling truth-telling notion when ε is small – as then the gains

¹⁴That is, $(v_\eta, v_1, \dots, v_K) \in \mathcal{V}$ implies $v_o \neq v_{o'}$ for any two distinct elements o and o' of \mathcal{O}' .

¹⁵The ε -Bayesian incentive compatibility truth-telling notion is weaker than other classical truth-telling notions, including: (i) “strategy-proofness,” which requires players find truth-telling optimal when they know each other’s types, and (ii) “Bayesian incentive compatibility,” which requires that ε -Bayesian incentive compatibility hold for $\varepsilon = 0$. It is, however, similar in spirit to other “approximate” truth-telling notions such as “strategy-proofness in the large” ([Azevedo & Budish, 2018](#)) and “asymptotic strategy-proofness” ([Liu & Pycia, 2016](#)).

Player	“Objects”			Probability	Player			Player	“Objects”				
	η	o	o'		1	2	3		η	o	o'		
1	θ_1	1	3	2	1/3	o	o'	η	1	θ'_1	1	5	3
2	θ_2	1	5	3	1/3	o'	o	η	2	θ'_2	1	5	3
3	θ_3	2nd	1st	3rd	1/6	o'	η	o	3	θ'_3	2nd	1st	3rd
					1/6	η	o'	o					

(a) Initial Types

(b) Deferred Acceptance with Uniform Tie-Breaking

(c) New Types

Table 1: Illustration

from manipulation are less than the (implicit) costs of manipulation. Hence, truth-telling is a Nash equilibrium of the reporting phase.

3. ILLUSTRATION

To illustrate our approach and give intuition to our core results, consider a simple example with three players, labeled 1 to 3, and two objects, labeled o and o' .¹⁶ Let the chance nature reveals a true value be $\alpha = \frac{1}{2}$, let \mathcal{V} have five elements – $(v_\eta, v_o, v_{o'}) = (1, 3, 2)$, $(v_\eta, v_o, v_{o'}) = (1, 5, 3)$, $(v_\eta, v_o, v_{o'}) = (2, 3, 0)$, $(v_\eta, v_o, v_{o'}) = (3, 4, 1)$, and $(v_\eta, v_o, v_{o'}) = (1, 6, 7)$ – and let each element be equally likely, i.e., $f_v = \frac{1}{5}$.

TYPE, PAYOFFS, AND AN INTERIM EFFICIENT ASSIGNMENT

Suppose the players' types are given by Table 3.1(a). So, players 1 and 2 learn their cardinal preferences and thus their payoffs, which are $u_1(y) = 1\mathbb{I}(y = \eta) + 3\mathbb{I}(y = o) + 2\mathbb{I}(y = o')$ and $u_2(y) = 1\mathbb{I}(y = \eta) + 5\mathbb{I}(y = o) + 3\mathbb{I}(y = o')$ respectively.¹⁷ Player 3 only learns her ordinal preference. Since $(v_\eta, v_o, v_{o'}) = (2, 3, 0)$ and $(v_\eta, v_o, v_{o'}) = (3, 4, 1)$ are the only true values consistent with 3's type and since all types are equally likely, her expected true values are

$$\frac{1}{5}(2, 3, 0) + \frac{1}{5}(3, 4, 1) = \left(\frac{5}{2}, \frac{7}{2}, \frac{1}{2}\right),$$

and her payoff is $u_3(y) = \frac{5}{2}\mathbb{I}(y = \eta) + \frac{7}{2}\mathbb{I}(y = o) + \frac{1}{2}\mathbb{I}(y = o')$.

¹⁶We refer to the objects as o and o' to avoid confusion with player names.

¹⁷We write $\mathbb{I}(\{\text{event}\})$ for the classic indicator function, which equals one when event is true and equals zero otherwise. In addition, we suppress the dependency of payoffs on type for expositional simplicity.

In the corresponding interim efficient assignment, player 1 gets o' , player 2 gets o , and player 3 gets η (i.e., 3 is left unassigned). Simply, its surplus maximizing to give objects to those who value them most, so: (i) the fact player 2 has the highest value for o (among all the players) implies it is efficient for her to get it and, given this, (ii) the facts players 1 and 3 have highest values for o' and η respectively (among all the remaining players) imply it is efficient they be assigned these items.

RECOMMENDATIONS OF CLASSIC MECHANISMS

Many classic mechanisms do not suggest the interim efficient assignment (when players report truthfully). For instance, Deferred Acceptance with Uniform Tie-Breaking, as well as Random Serial Dictatorship and Top Trading Cycles with Random Endowments, yield the mixed assignment in Table 3.1(b).¹⁸ This assignment is interim inefficient: its social surplus is $\frac{25}{3} \approx 8.33$, while the social surplus of the interim efficient assignment is $\frac{19}{2} = 9.5$, for an inefficiency of $\frac{7}{6}$.¹⁹

Intuitively, these mechanisms do not generate the interim efficient assignment because they “disregard cardinal information” and thus do not award objects to the players who value them most. To illustrate, consider Deferred Acceptance with Uniform Tie-Breaking. This mechanism consumes players reports and then, on behalf of each player, it “proposes” to the player’s most-preferred element in \mathcal{O}' . Each object then accepts one of its proposals, breaking ties at random; the object is then assigned to the player whose proposal it accepted. (Players who propose to η are always assigned to it.) Afterwards, the mechanism re-proposes on behalf of each player whose proposal was rejected

¹⁸ Deferred Acceptance traces its origins back to the classical “marriage market” of Gale & Shapley (1962), and finds more recent treatments in Abdulkadiroglu et al. (2011) and Roth & Sotomayor (1990); we describe it below. The Random Serial Dictatorship mechanism randomly orders players according to a uniform distribution, and then assigns the first her most-preferred element in \mathcal{O}' , the second her most preferred element among the remaining elements, and so on. Abdulkadiroglu & Sonmez (1998) establish that the mixed outcome of Random Serial Dictatorship is the same as under Top Trading Cycles with Random Endowments. This latter mechanism endows players randomly with objects or η and asks them to “point” to their most preferred objects, a cycle of the implied graph is then found and its trades implemented as assignments; the process repeats until there are no cycles left. Top Trading Cycles traces its origins back to Shapley & Scarf (1974), and finds modern treatments in Abdulkadiroglu & Sonmez (2003) and Roth et al. (2004).

¹⁹ While this inefficiency is small, inefficiencies in general can be quite large – e.g., if we multiply the true values by $\alpha > 1$, then the inefficiency is $\frac{7}{6}\alpha$.

to one of the remaining objects or η . The process concludes once all players are assigned to an object or η .

Since the three players initially propose to object o , there is a $\frac{1}{3}$ chance each gets it due to o 's uniform tie-breaking. Thus, the mixed assignment's outcome is inefficient 2/3-rds of the time.²⁰ The root of the inefficiency lies in the fact that o ignores players' expected true values when deciding which proposal to accept. Instead, if o broke ties in favor of higher values for itself, then (i) it would go to player 2 in the initial round and, in the second round, (ii) o' would go to player 1 and η would go to player 3. Thus, the interim efficient assignment would obtain with certainty.

This leads us to a classical conclusion of the random assignment literature: mechanisms that disregard cardinal information “can be” interim inefficient. In Proposition 1, we layout broad conditions which strengthen this conclusion and transform “can be” into “must be.” Roughly, these conditions stipulate that (i) there are a sufficient number of players and (ii) that \mathcal{V} contains two different, possible vectors of true values that induce the same ordinal preference. We thus obtain the following observation.

Observation 1. Mechanisms that disregard cardinal information are interim inefficient.

One might think that the inefficiency problem is limited only to mechanisms that disregard cardinal information. In Section 4, we show that this is not the case by establishing that classic mechanisms that use cardinal information, like Pseudo-markets, are also interim inefficient. The broad intuition is that mechanisms often use cardinal information in ways that are unrelated to social surplus maximization.

THE SIMPLE MECHANISM

Motivated by Observation 1, we introduce the “Simple Mechanism,” which is formally described in Section 5. In this mechanism, players submit either cardinal or ordinal reports. The mechanism uses these reports to construct an

²⁰The mixed assignment in Table 3.1(b) results from this tie-breaking. If player 1 wins o , then player 2 proposes to o' and player 3 proposes to η in the next round. Since both proposals are accepted, the assignment in the first row of Table 3.1(b) is obtained. If 2 wins o , then 1 proposes to o' and 3 proposes to η in the next round, so the assignment in the second row of Table 3.1(b) is obtained. If 3 wins o , then 1 and 2 both propose to o' in the next round, with the loser proposing to η in the third round, so the assignments in the last two rows of Table 3.1(b) are obtained. The probabilities of each pure assignment in Table 3.1(b) are product of the respective tie-breaking events.

estimate of each player's payoffs. It does this by treating a player's report as her type and computing her expected true values: specifically, (i) if a player submits a cardinal report r , then the mechanism estimates her expected true values are r , and if a player submits an ordinal report r , then the mechanism estimates her expected true values are $\sum_{v \in I(r)} (v f_v(v) / \sum_{v' \in I(r)} f_v(v'))$. The mechanism then computes the set of pure assignments σ that maximizes the sum of these estimated payoffs (i.e., maximize its estimate of social surplus) and it selects one of these assignments to implement with uniform probability.

To illustrate how the Simple Mechanism works, suppose the players truthfully report their types from Table 3.1(a). Then the mechanism estimates players 1 and 2 have payoffs of $\hat{u}_1(y) = 1\mathbb{I}(y = \eta) + 3\mathbb{I}(y = o) + 2\mathbb{I}(y = o')$ and $\hat{u}_2(y) = 1\mathbb{I}(y = \eta) + 5\mathbb{I}(y = o) + 3\mathbb{I}(y = o')$ respectively since they make cardinal reports. Since player 3 submits an ordinal report r_3 , the mechanism first determines that $I(r_3) = \{(2, 3, 0), (3, 4, 1)\}$ and estimates her expected true values are $(\frac{5}{2}, \frac{7}{2}, \frac{1}{2}) = \sum_{v \in I(\theta_3)} (v \frac{1}{5} / \sum_{v' \in I(\theta_3)} \frac{1}{5})$ and her payoff is $\hat{u}_3(y) = \frac{5}{2}\mathbb{I}(y = \eta) + \frac{7}{2}\mathbb{I}(y = o) + \frac{1}{2}\mathbb{I}(y = o')$. The mechanism then solves $\max_{\phi \in \Phi} \sum_{i \in \{1, 2, 3\}} \hat{u}_i(\phi(i))$ and concludes that the unique solution is for player 1 to get o' , player 2 to get o , and player 3 to get η . And it implements this solution with certainty.

The interim efficiency of this solution is by design: when players tell the truth, the Simple Mechanism correctly estimates their payoffs and thus maximizes social welfare when it maximizes the sum of their estimated payoffs. This leads to our second observation, which we formalize in Proposition 2.

Observation 2. The Simple Mechanism is interim efficient.

It follows that the Simple Mechanism is also interim and ex-post Pareto efficient, as well as individually rational. Nevertheless, the Simple Mechanism may be inefficient when players lie. To see this, suppose player 3 misreports her preference and states that o' is most preferred, o is second most preferred, and η is least preferred. Then, the mechanism estimates her payoff is $\hat{u}'_3(y) = 1\mathbb{I}(y = \eta) + 6\mathbb{I}(y = o) + 7\mathbb{I}(y = o')$, maximizes $\hat{u}_1 + \hat{u}_2 + \hat{u}'_3$, and concludes that there is a unique solution: assign η to player 1, assign o to player 2, and assign o' to player 3. This solution is clearly interim inefficient.

A positive feature of the Simple Mechanism is that it is "symmetric," i.e., it treats equals equally. To illustrate, suppose Table 3.1(c) gives the players' types, instead of Table 3.1(a); so, players 1 and 2 now have the same cardinal preference. Then, under truthful reporting, the Simple Mechanism computes

that there are two assignments in σ : (i) 1 gets o , 2 gets o' , and 3 gets η and (ii) 1 gets o' , 2 gets o , and 3 gets η . Since it mixes over these assignments with equal probability, players 1 and 2 both have a $\frac{1}{2}$ chance of getting o or o' and thus have the same payoff.

The intuition straightforward. Since players 1 and 2 have the same payoff function, it is always possible to have them exchange objects while retaining efficiency. Thus, for every assignment in σ where player 1 does better than player 2 by an amount δ , there is a second assignment in σ where 2 does better than 1 by δ . Since the Simple Mechanism randomizes uniformly over σ , both players have the chance of getting the better or worse assignment and thus have the same payoff. This leads us to our third observation, which we formalize in Proposition 2.

Observation 3. When two players have the same preferences and report truthfully, then they have the same payoff to the Simple Mechanism.

There is a question of whether players want to tell the truth to the Simple Mechanism. Unfortunately, truth-telling is neither strategy-proof nor Bayesian incentive compatible. To illustrate the lack of Bayesian incentive compatibility (and thus the lack of strategy-proofness),²¹ it helps to simplify the example by assuming that there are only two players, 1 and 2, that $\alpha = 1$, and that \mathcal{V} contains two equally likely elements $(v_\eta, v_o, v_{o'}) = (1, 5, 3)$ and $(v_\eta, v_o, v_{o'}) = (1, 3, 2)$, which we label A and B respectively. Focus on player 1 and suppose her type is B .

First, consider player 1's payoff to truth-telling, i.e., reporting she is a type B , when player 2 tells the truth. If player 2 is a type A , then the Simple Mechanism gives 2 object o and gives 1 object o' . If player 2 is type B , then the Simple Mechanism computes there are two assignments in σ : (i) 1 gets o and 2 gets o' and (ii) 1 gets o' and 2 gets o ; so, player 1 gets o half of the time and gets o' the other half of the time. Hence, player 1's payoff to truth-telling is $\frac{1}{2}2 + \frac{1}{2}(\frac{1}{2}3 + \frac{1}{2}2) = \frac{9}{4}$.

Second, consider player 1's payoff to lying, i.e., reporting she is a type A , when player 2 tells the truth. If player 2 is a type A , then the Simple Mechanism computes that there are two assignments in σ : (i) 1 gets o and 2 gets o' and (ii) 1 gets o' and 2 gets o ; so, player 1 gets o half of the time and gets o' the other

²¹ Strategy-proofness implies Bayesian incentive compatibility because the former requires truth-telling be optimal for each possible profile of types, whereas the latter integrates across type profiles.

half of the time. If player 2 is type B , then the Simple Mechanism concludes that it is best to give o to 1 and give o' to 2. Hence, player 1's payoff to lying is $\frac{1}{2}(\frac{1}{2}3 + \frac{1}{2}2) + \frac{1}{2}3 = \frac{11}{4}$.

Player 1's gain to lying is thus $\frac{1}{2} = \frac{11}{4} - \frac{9}{4}$. Intuitively, this gain is positive because the Simple Mechanism awards objects to those who it thinks value them most. Thus, if player 1 does not report the highest value for o' , then she will not get it with positive probability. Since player 1 knows this and knows she has the lowest value for o' , she finds it best to lie and misreport her type.²²

Importantly, player 1's gain to lying decreases as the number of truth-telling players $N - 1$ grows. To see this, let A_l be the event that there are l players (aside from player 1), who report type A , i.e., $A_l = \{\theta_{-1} | \theta_{-1} \text{ specifies } l \text{ players have type } A\}$. We can then write 1's payoff to reporting r_1 as

$$U_1^{M_S}(r_1|B) = \sum_{l=0}^{N-1} \sum_{\theta_{-1} \in A_l} u_1(M_S(r_1, \theta_{-1})|B) \Pr(\theta_{-1}),$$

where M_S denotes random assignment selected by the Simple Mechanism. It is straightforward to compute the value of $u_1(M_S(r_1, \theta_{-1})|B)$ on each set A_l :

- For $l = 0$ and all θ_{-1} in A_l , we have

$$u_1(M_S(r_1, \theta_{-1})|B) = \begin{cases} 3 & \text{if } r_1 = A \\ \frac{3}{N} + \frac{2}{N} + 1(1 - \frac{2}{N}) = 1 + \frac{3}{N} & \text{if } r_1 = B. \end{cases}$$

Since type A s value both o and o' strictly more than type B s, the Simple Mechanism gives o and o' to players who report type A in order to maximize its estimate of social surplus. Thus, (i) if player 1 reports A , then she gets o since there are no other type A players, and (ii) if she reports B , then she gets o' with probability $\frac{1}{N}$, o with probability $\frac{1}{N}$, and η with probability $(1 - \frac{2}{N})$.²³

²² Ehlers et al. (2014) leverage an analogous intuition to show that every mechanism that awards objects based on cardinal information is not Bayesian incentive compatible in finite assignment environments.

²³ These probabilities come from a counting exercise. There are $\binom{N}{1} \binom{N-1}{1}$ ways to distribute the two objects to the N players who report B . Thus, if player 1 has o (or o'), there are $\binom{N-1}{1}$ ways to distribute o' (o) to the other players. Since all assignments in σ are equally likely, the probability player 1 gets o (o') is $\binom{N-1}{1} / (\binom{N}{1} \binom{N-1}{1}) = \frac{1}{N}$. Since player 1 gets η if she does not get o or o' , the probability of η is $1 - \frac{2}{N}$. (The counting exercises for all subsequent bullets are analogous and omitted.)

- For $l = 1$ and all θ_{-1} in A_l , we have

$$u_1(M_S(r_1, \theta_{-1})|B) = \begin{cases} \frac{3}{2} + \frac{2}{2} = \frac{5}{2} & \text{if } r_1 = A \\ \frac{2}{N-1} + 1(1 - \frac{1}{N-1}) = 1 + \frac{1}{N-1} & \text{if } r_1 = B. \end{cases}$$

If $r_1 = A$, then the Simple Mechanism assigns both objects to the two A reporters, implying that player 1 has an equal chance of getting either object. If $r_1 = B$, then because the player who truthfully reports A always gets o , the chance player 1 gets o' is $\frac{1}{N-1}$.

- For $l \geq 2$ and all θ_{-1} in A_l , we have

$$u_1(M_S(r_1, \theta_{-1})|B) = \begin{cases} \frac{1}{l+1}3 + \frac{1}{l+1}2 + (1 - \frac{2}{l+1})1 = 1 + \frac{1}{l+1}3 & \text{if } r_1 = A \\ 1 & \text{if } r_1 = B. \end{cases}$$

If $r_1 = A$, then player 1 gets o with probability $\frac{1}{l+1}$, o' with probability $\frac{1}{l+1}$, and η with probability $(1 - \frac{2}{l+1})$. If $r_1 = B$, then player 1 gets η because o and o' go to the A reporters.

Hence,

$$U_1^{M_S}(r_1|B) = \begin{cases} 3\Pr(A_0) + \frac{5}{2}\Pr(A_1) + \sum_{l=2}^{N-1}(1 + \frac{3}{l+1})\Pr(A_l) & \text{if } r_1 = A \\ (1 + \frac{3}{N})\Pr(A_0) + (1 + \frac{1}{N-1})\Pr(A_1) + \sum_{l=2}^{N-1}\Pr(A_l) & \text{if } r_1 = B, \end{cases} \quad (3)$$

where, due to the independence of players' types, $\Pr(A_l) = \sum_{\theta_{-1} \in A_l} \Pr(\theta_{-1}) = \binom{N-1}{l} \frac{1}{2}^{N-1}$.²⁴ Observe that $U_1^{M_S}(r_1|B)$ is greater than one and decreases to one as N grows large.²⁵ Thus, for each $\varepsilon > 0$, there is a finite $\bar{N} > 0$ such that $1 < U_1^{M_S}(A|B) \leq 1 + \varepsilon$ and $1 < U_1^{M_S}(B|B) \leq 1 + \varepsilon$ when there are at least \bar{N}

²⁴There are $\binom{N-1}{l}$ elements in A_l and the probability of each is $(\frac{1}{2})^{N-1}$ since types are independent, $\alpha = 1$, and all both types are equally likely.

²⁵We establish that $\lim_{N \rightarrow \infty} U_1^{M_S}(A|B) = 1$ since the argument that $\lim_{N \rightarrow \infty} U_1^{M_S}(B|B) = 1$ is analogous. Since $\lim_{N \rightarrow \infty} \Pr(A_l) = 0$, it is evident that the first two terms on the right-hand-side of equation (3) go to zero. We use a squeeze argument to show that the last term goes to one. Let $\delta > 0$. Since $\frac{3}{l+1}$ is strictly decreasing, there is a fixed N' such that (i) $\frac{3}{l+1} > \delta$ for all $l < N'$ and (ii) $\frac{3}{l+1} \leq \delta$ for all $l \geq N'$. Thus, for large N , we have $1 \leq \sum_{l=2}^{N-1}(1 + \frac{3}{l+1})\Pr(A_l) \leq \sum_{l=2}^{N'-1}(1 + \frac{3}{l+1})\Pr(A_l) + (1 + \delta)\sum_{l=N'}^{N-1}\Pr(A_l)$. Since $\sum_{l=N'}^{N-1}\Pr(A_l) = (1 - \sum_{l=0}^{N'-1}\Pr(A_l))$, taking limits yields $1 \leq \lim_{N \rightarrow \infty} \sum_{l=2}^{N-1}(1 + \frac{3}{l+1})\Pr(A_l) \leq 1 + \delta$. Letting $\delta \rightarrow 0$ gives the desired result.

players. In other words, player 1 never gains more than ε by lying when there are at least \bar{N} players.²⁶

The economic intuition is simple: “competition” for the objects increases as the number of truthful type A players rises (due to the growth of N). This competition is driven by the Simple Mechanism’s randomization, which ensures that each A reporter has an equal chance of getting either object. Thus, the probability that player 1 can obtain either object by lying or telling the truth falls, limiting her ability to manipulate the mechanism. Since player 1 is symmetric to all type B players and since no type A player has incentive to lie, we obtain the following observation, which we formalize in Proposition 3.

Observation 4. For every $\varepsilon > 0$, the Simple Mechanism is ε -Bayesian incentive compatible for large N .

The Simple Mechanism thus has good truth-telling properties when the number of players is large. This, along with interim efficiency and symmetry, suggest that the Simple Mechanism may be a good mechanism to employ in situations there is substantial demand for the objects – as might be the case in the assignment of prime university housing or courses, in the awarding of grants or fellowships, in the allocation of slots in programs for free/subsidized day-care or housing, in the apportionment of prime employee parking spots or offices, or even in the distribution of kidneys and other organs.

We have held the set of objects fixed while increasing the number of players. While there are situations where the supply of objects is fixed, it is often natural to think that it increases with the number of players – e.g., colleges build dorm rooms and add courses in response to rising enrollments and local governments expand free day-care in response to growing populations. Fortunately, this assumption is for expositional simplicity. As we discuss in Proposition 4, the set of objects may increase with the number of players via replication, provided it does so slowly enough to maintain competitive pressure and thus approximate incentive compatibility.

²⁶In practice, N does not have to be very large for the gain to lying to be small – e.g., computation gives that player 1’s gain to lying is 0.12 when there are 50 players and 0.06 when there are 100 players.

4. CARDINAL INFORMATION AND INEFFICIENCY

Intuitively, mechanisms that “disregard cardinal information” can be interim inefficient. Our main goal, in this section, is to strengthen this intuition and show that such mechanisms “must be” interim inefficient under broad conditions. We also discuss the efficiency of classical mechanisms that make use of cardinal information.

DISREGARDING CARDINAL INFORMATION

We have in mind that a mechanism “disregards cardinal information” if it only uses the ordinal information in players’ reports. To make this notion precise, for each $\theta \in \Theta$, let $c(\theta)$ give the order that is consistent with θ . Thus, (i) $c(\theta) = \theta$ if $\theta \in \mathcal{B}'$ or (ii) $c(\theta) = \preceq$ if $\theta \in \mathcal{V}$ where \preceq is the element of \mathcal{B}' that is consistent with θ . For a type profile $\theta = (\theta_1, \dots, \theta_N)$ or a report profile $r = (r_1, \dots, r_N)$, we write $c(\theta)$ and $c(r)$ for $(c(\theta_1), \dots, c(\theta_N))$ and $(c(r_1), \dots, c(r_N))$ respectively for players ordinal preferences and ordinal reports. We say that a mechanism M **disregards cardinal information** if $M(r) = M(c(r))$ for all $r \in \Theta^N$. Many classic mechanisms disregard cardinal information since they only use players’ rankings of objects when they run – e.g., Deferred Acceptance, Probabilistic Serial,²⁷ Random Serial Dictatorship, and Top Trading Cycles.

To establish the interim inefficiency of mechanisms that disregard cardinal information, we place a regularity condition on \mathcal{O} and \mathcal{V} .

Assumption 1. There exist two objects o and o' in \mathcal{O} and two vectors $\check{v} = (\check{v}_\eta, \dots, \check{v}_o, \dots, \check{v}_{o'}, \dots, \check{v}_K)$ and $\check{v}' = (\check{v}'_\eta, \dots, \check{v}'_o, \dots, \check{v}'_{o'}, \dots, \check{v}'_K)$ in \mathcal{V} such that:

- (i) Both vectors induce the same ordinal preferences, i.e., $c(\check{v}) = c(\check{v}')$.
- (ii) Object o' is strictly preferred to o and o is strictly preferred to η , i.e., $\check{v}_{o'} > \check{v}_o > \check{v}_\eta$ and $\check{v}'_{o'} > \check{v}'_o > \check{v}'_\eta$.
- (iii) The “marginal” value to giving up o to get o' is larger in \check{v}' than \check{v} , i.e., $\check{v}'_{o'} - \check{v}'_o > \check{v}_{o'} - \check{v}_o$.

Assumption 1 is a mild condition that is satisfied, for instance, in the Illustration or, more generally, when $\mathcal{V} = \{v_0, v_1, \dots, v_L\}^{K+1}$ with $v_0 < v_1 < \dots < v_L$,

²⁷The Probabilistic Serial mechanism was introduced by [Bogomolnaia & Moulin \(2001\)](#). In this mechanism, players consume their most preferred object at a some “eating speed.” Once the objects are consumed, players move onto their second most preferred objects, provided they are available, and so on until all objects are consumed. The proportion of each object eaten by player i then corresponds to the probability that i gets the object in the assignment.

$K \geq 2$, and $L \geq 3$.²⁸

Proposition 1. Disregarding Cardinal Information and Inefficiency.

Let Assumption 1 hold, let $\alpha > 0$, and let M be a mechanism that disregards cardinal information. Then, there is a finite $\bar{N} > 0$ such that M is interim inefficient when there are at least \bar{N} players.

The proposition implies, for instance, that Deferred Acceptance, Probabilistic Serial, Random Serial Dictatorship, and Top Trading Cycles all produce assignments with social surplus that is *strictly* less than what could be achieved for at least one type profile. The size of the inefficiencies may be economically substantial, especially when the values in \mathcal{V} are large or “spread out.”

The intuition for the result mirrors the intuition behind Observation 1; we exploit it in the proof by initially endowing all players with the same type and selecting a player who gets o under M . We then increase this player’s value to o' so that it becomes efficient for her alone to get it. Since this cardinal perturbation leaves ordinal types unaltered, M still assigns o to our player and is thus interim inefficient. Assumption 1 enables this preference alteration.

The lower bound on the number of players rules out corner cases where interim efficiency may obtain even for mechanisms that disregard cardinal information – e.g., when there is a single player, Randomized Serial Dictatorship always allows her to obtain her most preferred object, an interim efficient outcome. The restriction that $\alpha > 0$ ensures that some players have cardinal preferences; if $\alpha = 0$, then all preferences are ordinal and the proposition is vacuous.

MAKING USE OF CARDINAL INFORMATION

In light of Proposition 1, it is natural to wonder if mechanisms that “use cardinal information” are interim efficient. The answer, unfortunately, is “not always:” some, like the Simple Mechanism, are efficient, while other classic mechanisms, like “Nash-Boston” or Pseudo-markets, are not.²⁹

Nash-Boston Mechanism. While the Boston mechanism only solicits players’ strict ordinal preferences and so is interim inefficient, the equilibrium of its reporting game depends on players best responses, which themselves

²⁸ When $K \geq 2$ and $L \geq 3$, there are value vectors $(v_0, v_0, \dots, v_0, v_1, v_2)$ and $(v_0, v_0, \dots, v_0, v_1, v_3)$, which are easily seen to meet (i), (ii), and (iii) of Assumption 1 because $v_3 > v_2 > v_1$.

²⁹ Interestingly, when there is a continuum of players and a mass of objects, both Nash-Boston and Pseudo-markets are interim efficient for a broad class of environments – see [Hafalir & Miralles \(2014\)](#).

depend on their preference intensities; see, for instance, [Abdulkadiroglu et al. \(2011\)](#) and [Miralles \(2008\)](#). It is thus natural to ask whether this Nash equilibrium – which we refer to as the “Nash-Boston” mechanism since it maps types into allocations – constitutes an interim efficient mechanism. As the next example shows, the answer is no.

Example 1. Nash-Boston and Interim Inefficiency.

Let there be two players, 1 and 2, and two objects, o and o' . Let $\alpha = 1$ and let \mathcal{V} contain two equally likely elements $(v_\eta, v_o, v_{o'}) = (0, 10, -1)$ and $(v_\eta, v_o, v_{o'}) = (0, 9, 8)$, which we label A and B respectively. It is clear that if both players are type A (B), then the two interim efficient pure assignments give o to one player and η to the other player (give o to one player and o' to the other player). However, if there is one of each type, then the interim efficient assignment gives o to the type A player and o' to the type B player.

Computation shows that, in the (unique) symmetric equilibrium of the Boston Mechanism’s reporting game, both types of players truthfully report their ordinal preferences.³⁰ This equilibrium has the following outcomes. If both players are type A (B), one is randomly assigned o and the other is assigned η (assigned o and the other is assigned o'). If one player is type A and the other is type B , one is randomly assigned o and the other is assigned either η if type A or o' if type B . In all outcomes, both players have an equal probability of being assigned o .

The equilibrium type-assignment mapping, i.e., the Nash Boston mechanism, produces a mixed assignment that randomizes over the two interim efficient pure assignments when both players are type A (B) and so itself is interim efficient. However, when one player is type A and the other type B , then the mechanism produces a mixed assignment that gives o to the type B player with non-zero probability and so is interim inefficient. It follows that the Nash-Boston mechanism is interim inefficient: intuitively, while the equilibrium consumes preference intensities via players’ best responses, this information is used to form a mutual best response and not to cooperatively maximize social surplus. \triangle

Pseudo-Market Mechanisms. A Pseudo-market mechanism assigns goods according to the Walrasian equilibrium of an artificial economy that is constructed from players’ reported preferences; see, for instance, [Hylland &](#)

³⁰We defer the details of the Boston mechanism, its reporting game, and the equilibrium to the Appendix for expositional simplicity.

Zeckhauser (1979).³¹ Specifically, the mechanism endows each player with an artificial budget, which is specified in the mechanism's description, and then maximizes her reported utility by buying probability shares of the objects subject to this constraint, while taking prices as given. It then solves for the Walrasian equilibrium and uses the equilibrium probability shares as the basis for the assignment. Since the Pseudo-market uses cardinal information in the purchase of probability shares, one might think that its equilibrium is interim efficient. The next example shows that this is not the case.

Example 2. A Pseudo-Market and Interim Inefficiency.

Let there be three players, 1, 2, and 3, and two objects, o and o' . Let $\alpha = 1$ and let \mathcal{V} contains three equally likely elements $(v_\eta, v_o, v_{o'}) = (0, 4, 2)$, $(v_\eta, v_o, v_{o'}) = (0, 3, 2)$, and $(v_\eta, v_o, v_{o'}) = (0, 2, 4)$, which we label A , B , and C respectively.

To begin, suppose that nature draws type A for player 1, type B for player 2, and type C for player 3, i.e., types are $\theta = (A, B, C)$. Then, it is uniquely interim efficient for o go to player 1, for η to go to player 2, and for o' go to player 3.

Consider a Pseudo-market mechanism M , where (i) players 1 and 3 have budgets of one and (ii) player 2 has a budget of zero. Suppose the prices for o and o' , denoted p_o and $p_{o'}$, are $p_o = p_{o'} = 1$. On behalf of player 1, the mechanism solves $\max_{(x_o, x_{o'})} 4x_o + 2x_{o'}$ such that $p_o x_o + p_{o'} x_{o'} \leq 1$, where x_o and $x_{o'}$ are the probability shares of o and o' respectively. It concludes that the solution is $(x_o^*, x_{o'}^*) = (1, 0)$ and thus purchases a unit of probability shares of o for player 1. The mechanism solves an analogous problem for player 3 and purchases a unit of probability shares of o' for player 3. Since player 2 has a budget of zero, the mechanism does not purchase any probability shares on her behalf. Since there is a unit supply of probability shares of each object, markets clear and the prices and shares constitute the Walrasian equilibrium. Since the equilibrium probability shares determine the assignment, with a unit share indicating a 100% chance of ownership, $M(\theta)$ assigns o to player 1, η to player 2, and o' to player 3.

It turns out that M is interim inefficient at any type profile where it is efficient for player 2 to receive an object. To illustrate, suppose nature draws B for player 1, A for player 2, and C for player 3, i.e., types are $\theta' = (B, A, C)$

³¹ More recent treatments of Pseudo-markets and their derivative mechanisms can be found, for instance, in Budish (2011), Budish et al. (2013), and He et al. (2018).

instead of θ . Then, it is uniquely interim efficient for η to go to player 1, for o to go to player 2, and for o' to go to player 3. Yet, this assignment cannot be the outcome of $M(\theta')$ because budgets are pre-specified: simply, player 2 has a budget of zero, so if the mechanism is to purchase any probability shares of o on her behalf, it must be that $p_o = 0$. Yet, if $p_o = 0$, the mechanism would purchase infinite quantities of probability shares of o for all three players and markets would not clear. Hence, $M(\theta')$ is interim inefficient.

More generally, if it is interim efficient for player i to receive an object at one type profile and to receive η under another type profile, then no Pseudo-market mechanism is interim efficient because efficiency requires i have a non-zero budget under the former profile and a zero budget under the latter profile, an impossibility. It is thus the fixed nature of budgets that drives the Walrasian equilibrium away from an interim efficient assignment, despite the fact the equilibrium consumes cardinal information. \triangle

In light of this example, one might consider developing an “alternative” Pseudo-market mechanism M by varying a Pseudo-market’s budgets in response to players’ reports and then allocating objects according to one of its equilibria. [Miralles & Pycia \(2014, Theorem 1\)](#) show that every Pareto efficient assignment corresponds to some Walrasian equilibrium of some Pseudo-market. Thus, for every profile of types θ , if one were to set the budgets in the right way, one could ensure that one of the resulting Walrasian equilibria was interim efficient and have $M(\theta)$ allocate objects according to this equilibrium. However, such a mechanism would be computationally difficult because there are no known, exact methods to compute either the necessary budgets or the equilibrium prices. We thus develop an alternative, interim efficient mechanism that is based on a linear programming problem.

5. THE SIMPLE MECHANISM

In the Illustration, we informally introduced the Simple Mechanism as a reaction to the interim inefficiency of other classic mechanisms and informally developed at its properties. Our goal, in this section, is to formally describe this mechanism and its core properties.

The Simple Mechanism.³²

³²This mechanism benefited greatly from discussions with Marek Pycia.

Suppose the players report $r = (r_1, \dots, r_N) \in \Theta^N$. Then, the mechanism returns the mixed assignment $M_S(r)$, which is constructed as follows. First, for each player i , construct an estimate \hat{v}_i of i 's expected true values by setting

$$\hat{v}_i = (\hat{v}_{\eta i}, \hat{v}_{1i}, \dots, \hat{v}_{Ki}) = \begin{cases} r_i & \text{if } r_i \in \mathcal{V} \\ \sum_{v \in I(r_i)} v \frac{f_v(v)}{\sum_{v \in I(r_i)} f_v(v)} & \text{if } r_i \in \mathcal{B}', \end{cases} \quad (4)$$

and represent \hat{v}_i as a function \hat{u}_i by setting $\hat{u}_i(o) = \hat{v}_{oi}$ for each $o \in \mathcal{O}'$. Second, compute the set

$$\sigma = \arg \max_{\phi \in \Phi} \sum_{i \in \mathcal{N}} \hat{u}_i(\phi(i)). \quad (5)$$

Third, set $M_S(r)$ to assign weight $1/|\sigma|$ to each pure assignment in σ .³³○

In words, the Simple Mechanism takes in players' reports, estimates their expected true values,³⁴ computes the set of pure assignments that maximize the sum of these estimated values, and then selects one of these assignments to implement with uniform probability.

The next two propositions formally establish the properties of the Simple Mechanism that we outlined in Observations 2 to 4: interim efficiency, symmetry, and ε -Bayesian incentive compatibility. All proofs are provided in the Appendix.

We say that a mechanism M is **symmetric** if, when players i and k make the same report $r_i = r_k$ and when they have the same type $\theta_i = \theta_k$, they have the same payoff to the (mixed) assignment $M(r)$, where $r = (r_1, \dots, r_i, \dots, r_k, \dots, r_N) \in \Theta^N$.

Proposition 2. The Simple Mechanism is Interim Efficient and Symmetric. *The Simple Mechanism is interim efficient and symmetric.*

Interim efficiency implies that, under truthful reporting, the Simple Mechanism is interim Pareto efficient, ex-post Pareto efficient, and ordinally efficient

³³ Since Φ is finite, σ non-empty and finite.

³⁴ The Simple Mechanism requires knowledge of the structure of preferences to carry out its estimation; we examine the robustness of the mechanism with respect to this requirement in the Online Appendix. Also, a referee insightfully observed this requirement “shifts the [computational] burden” from players to the mechanism. From a mathematical perspective, one could shift it back by requiring players submit estimates of their expected true values to a mechanism, which assigns objects to maximize the sum of reported values. Provided this mechanism randomizes uniformly over the set of maximizing assignments, it has similar properties to the Simple Mechanism.

(when true values are strict); it thus is also individually rational.³⁵ Symmetry implies that the Simple Mechanism is “fair.” The result’s intuition mirrors the intuitions for Observations 2 and 3, and the proof amounts to a formalization of these intuitions.

Our incentive compatibility result relies on \mathcal{V} having particular type of structure. Specifically, we say that \mathcal{V} is a **product of finite grids** if $\mathcal{V} = \times_{o \in \mathcal{O}'} \mathcal{V}_o$, where $\mathcal{V}_o = \{v_1^o, v_2^o, \dots, v_{L_o-1}^o, v_{L_o}^o\}$ for some $0 \leq v_1^o < v_2^o < \dots < v_{L_o}^o \leq 1$ and $L_o \geq 1$ for each $o \in \mathcal{O}'$. For instance, if $\mathcal{V} = \{0, \frac{1}{3}, \frac{1}{2}, \frac{2}{3}, 1\}^{K+1}$ or if $\mathcal{V} = \{0\} \times \{0, \frac{1}{2}, 1\}^K$, then it is a product of finite grids. This mild restriction on \mathcal{V} ensures that, with positive probability, we can find a player who has a higher value for one object and a lower value for another object. The restriction of \mathcal{V} to the unit hyper-cube is without loss because we can always rescale the upper and lower bounds.

Proposition 3. The Simple Mechanism is ε -Bayesian Incentive Compatible. *Let \mathcal{V} be a product of finite grids and let $\alpha > 0$. Then, for each $\varepsilon > 0$, there is a finite $\bar{N} > 0$ such that the Simple Mechanism is ε -Bayesian incentive compatible when there are at least \bar{N} players.*

Since strategic manipulation is inherently costly, this proposition implies that the Simple Mechanism has good truth-telling properties when the number of players is large. The economic intuition is akin to the economic intuition behind Observation 3: as more players are added, “competition” for objects increases, which drives the probability a player can acquire any object, besides η , down to zero. This limits the opportunity for strategic manipulation and leads to ε -Bayesian incentive compatibility. Increasing competition itself is driven by the mechanism’s randomization and the finite grid assumption, which ensures that, for each object, players with high values for the object enter and “drive up the demand” as N grows. The technical approach of the proof mirrors the argument used in Section 3, but differs at key points due to the generality of the environment.³⁶

³⁵ It is evident that there are conditions under which the Simple Mechanism is “strictly” individually rational, i.e., all players do strictly better by participating in it than by remaining unassigned. For example, this is the case when (i) η is the least preferred option for all players and when (ii) there is at least one object that all players strictly prefer to η . It is also the case when (i’) all players have the same most preferred object and (ii’) value this object strictly more than η .

³⁶ The proof suggests a more general and abstract sufficient condition for ε -Bayesian incentive compatibility than finite grids, which we discuss in the Appendix.

A key assumption of the intuition (and of the proof) is that the set of objects \mathcal{O} is held constant while N grows. To dispense with this assumption, suppose \mathcal{O} grows via replication. A **linking function** is a $\psi : \mathbb{N}_+ \rightarrow \mathbb{N}_+$ that is monotone and has $\lim_{n \rightarrow \infty} \psi(n) = \infty$. We have in mind (i) that $\psi(N)$ gives the number of copies of each object and (ii) that players treat copies of objects as indistinguishable; we offer a formal treatment in the Appendix.

Proposition 4. Incentive Compatibility with Replication.

Let \mathcal{V} be a product of finite grids and let $\alpha > 0$. For each $\varepsilon > 0$:

- (i) When the linking function ψ is arbitrary, there may not exist a symmetric, interim efficient, and ε -Bayesian compatible mechanism.*
- (ii) There is a linking function $\bar{\psi}$ and an integer $\bar{N} > 0$ such that the Simple Mechanism is ε -Bayesian incentive compatible when there are at least \bar{N} players.*

The first part of the proposition is based on an example given in the Appendix. The example has $\psi(N) = \lfloor N/2 \rfloor$ copies of a single object and two equally likely types of players, A and B . The type A s value the object at 1, the type B s value it at $\frac{1}{2}$, and both value being unassigned at 0. Since the expected number of type A s equals the number of objects, any interim efficient and symmetric mechanism (including the Simple Mechanism) essentially assign all copies of the object to the type A s, leaving the type B s unassigned. A “back of the envelope calculation” thus suggests that if a type B lies, then she obtains the object with probability $\frac{N/2}{N/2+1} = \frac{1}{1+2/N}$ (per symmetry), which converges to 1 as $N \rightarrow \infty$; whereas she gets nothing if she tells the truth. While actual mathematics are more involved, this simple computation lays bare the intuition and establishes that every type B gains from lying, even as N grows; economically speaking, it shows that replication can negate competition.

The second part of the proposition is a consequence of Proposition 3. To see this, suppose N starts growing from one and \mathcal{O} starts off with one copy of each object. We wish to know at which value of N we can add a second copy of each object without compromising ε -Bayesian incentive compatibility. To find this value, assume for a moment that the number of copies does *not* change with N . Then, Proposition 3 tells us that there are thresholds N_1 and N_2 above which the Simple Mechanism is ε -Bayesian incentive compatible when \mathcal{O} contains, respectively, one copy of each object or two copies of each object. Thus, when N is growing and \mathcal{O} starts off with one copy of each object, (i) ε -Bayesian incentive compatible holds once N exceeds N_1 and

(ii) we can add a second copy of each object once N grows larger than N_2 without losing approximate incentive compatibility. Continuing this exercise defines thresholds N_3, N_4, \dots for three, four, or more copies and so implicitly defines the required linking function. Broadly, any linking function arrived at in this way has the number of copies increase sufficiently slowly with N so that competition increases (due to the use of Proposition 3).³⁷ (We provide a second example in the Appendix where we slow the replication in the prior example from $\lfloor N/2 \rfloor$ down to $\lfloor \ln(N)/2 \rfloor$ and establish that this is sufficient to ensure the Simple Mechanism is ε -Bayesian incentive compatible.)

A. APPENDIX: PROOFS AND ADDITIONAL EXAMPLES

This Appendix collects the proofs of our results, as well as additional examples, which are presented in the order that they are referenced in the main text.

Proof of Lemma 1. Obvious and omitted. \square

Proof of Lemma 2. Obvious and omitted. \square

Proof of Proposition 1. Let $\bar{N} = |\{\hat{o} \in \mathcal{O} | \check{v}_{\hat{o}} > \check{v}_{\eta}\}|$ and let $N > \bar{N}$. Our goal is to show that, M generates an interim inefficient assignment for some type profile.

To begin, suppose the type profile is $\theta = (\check{v}, \dots, \check{v})$ and let ϕ be a pure assignment in the support of $M(\theta)$. There are two cases to consider: (i) either o or o' is unassigned by ϕ or (ii) both objects are assigned by ϕ .

For case (i), let $\tilde{o} \in \{o, o'\}$ denote an unassigned object. Since (i) there are $|\{\hat{o} \in \mathcal{O} | \check{v}_{\hat{o}} > \check{v}_{\eta}\}|$ objects that all players value more than being unassigned and (ii) since $N > |\{\hat{o} \in \mathcal{O} | \check{v}_{\hat{o}} > \check{v}_{\eta}\}|$, at least one player is either assigned η or is assigned some other object o'' that she values no more than η . Let i denote such a player. Construct a new assignment ϕ' by giving \tilde{o} to player i , while leaving everyone else's assignment alone – i.e., set $\phi'(i) = \tilde{o}$ and $\phi'(k) = \phi(k)$ for all players $k \neq i$. Observe that $\sum_{l \in \mathcal{N}} u_l(\phi' | \theta_l) - \sum_{l \in \mathcal{N}} u_l(\phi | \theta_l) = u_i(\phi' | \theta_i) - u_i(\phi | \theta_i) = \check{v}_{\tilde{o}} - u_i(\phi | \theta_i) > 0$ since \tilde{o} is strictly preferred to η . Thus, ϕ is interim inefficient. Since a mixed assignment is interim efficient if and

³⁷ It is also possible to obtain approximate incentive compatibility by having the number of copies increase sufficiently quickly with N such that scarcity “evaporates” and everyone gets their most preferred object – e.g., set $\psi(N) \geq N$ for all N . However, such linking functions are economically uninteresting.

only if each pure assignment in its support is interim efficient, it follows that $M(\theta)$ is also interim inefficient.

For case (ii), let i be the player who gets o and let k be the player who gets o' . Consider a change in i 's type from \check{v} to \check{v}' . Formally, consider a new type profile $\theta' = (\theta'_1, \dots, \theta'_{i-1}, \theta'_i, \theta'_{i+1}, \dots, \theta'_N)$, where $\theta'_i = \check{v}'$ and $\theta'_l = \check{v}$ for all players $l \neq i$; this type profile occurs with positive probability per Assumption 1 and the fact $\alpha > 0$.

Since (i) M disregards cardinal information and (ii) $c(\theta') = c(\theta)$, we have $M(\theta') = M(c(\theta')) = M(c(\theta)) = M(\theta)$. Thus, ϕ is in the support of $M(\theta')$. Construct a new assignment ϕ' from ϕ by swapping i and k 's assignments, while leave all other players' assignments alone – i.e., set $\phi'(i) = o''$, $\phi'(k) = o'$, and $\phi'(l) = \phi(l)$ for all players $l \neq i, k$. We have

$$\begin{aligned} \sum_{l \in \mathcal{N}} u_l(\phi' | \theta'_l) - \sum_{l \in \mathcal{N}} u_l(\phi | \theta'_l) &= u_i(\phi' | \theta'_i) - u_i(\phi | \theta'_i) + u_k(\phi' | \theta'_k) - u_k(\phi | \theta'_k) \\ &= \check{v}'_{o'} - \check{v}'_{o''} - (\check{v}_{o'} - \check{v}_{o''}) > 0, \end{aligned}$$

where the strict inequality follows from Assumption 1. That is, ϕ is interim inefficient, implying $M(\theta')$ is interim inefficient. \square

Example 1 Details. The Nash Boston Symmetric Equilibrium.

In the Boston Mechanism's reporting game, players submit (strict) ordinal reports.³⁸ The mechanism then allocates objects by sequentially considering all players' first choices (in the first round), second choices (in the second round), and so on; it stops once each player is assigned an object or η . In the k -th round, the mechanism first "pins" the name of each player (who was not assigned to an object or η in a previous round) to her k -th choice from \mathcal{O}' , provided this choice is available (i.e., was not given to another player in a previous round); η is always available. If an object has one player's name pinned to it, then it is assigned to that player. However, if an object has two or more players' names pinned to it, then a (uniform) random lottery is held and the object is assigned to the winner. The item η is assigned to all players whose names are pinned to it. Players who lose their lotteries, players who have k -th choices that are unavailable, and objects that had no names pinned to them all move forwards into the next round.

As is common in the literature on the Boston mechanism (e.g., [Abdulkadiroglu et al., 2011](#)), we look for a symmetric Nash equilibrium. Observe that

³⁸ [Abdulkadiroglu & Sonmez \(2003\)](#) provide a detailed description of the Boston mechanism, which we have adapted for our single-unit-demand environment.

if a player's k -th choice is η and they enter the k -th round, then the mechanism assigns them η for certain. Thus, it is dominated for a type A to report o below η or to report o' above η – so, her dominant strategy is $o \succ \eta \succ o'$ – and it is dominated for a type B to report o or o' below η – so, she has two possible strategies: $s_1 := o \succ o' \succ \eta$ or $s_2 := o' \succ o \succ \eta$.³⁹

If B plays s_1 , then with probability $\frac{1}{2}$ she faces a type A . Since both players list o first, the mechanism conducts a fair lottery for it in round one; if the type B wins she gets it and if she loses then she gets o' in the next round. Hence, her payoff to facing a type A is $\frac{9}{2} + \frac{8}{2}$. A similar argument gives that her payoff to facing a type B is $\frac{9}{2} + \frac{8}{2}$ because of symmetry. Thus, her payoff to s_1 is $\frac{1}{2}(\frac{9}{2} + \frac{8}{2}) + \frac{1}{2}(\frac{9}{2} + \frac{8}{2})$. If B plays s_2 , then her payoff is $\frac{1}{2}8 + \frac{1}{2}(\frac{9}{2} + \frac{8}{2})$ via analogous logic. Thus, a type B 's best response is s_1 .

It follows that the symmetric Nash equilibrium is $o \succ \eta \succ o'$ if a type A and $o \succ o' \succ \eta$ if a type B . Computation then gives the mixed assignment for each type-profile reported in the main text. \triangle

Proof of Proposition 2. Suppose the players types are $\theta = (\theta_1, \dots, \theta_N)$, so their expected true values are $\{v_i^\dagger\}_{i \in \mathcal{N}}$ per equation (2) and their payoffs are $\{u_i\}_{i \in \mathcal{N}}$. We first wish to show that $M_S(\theta)$ is interim efficient. When the Simple Mechanism receives report θ , it computes estimates of the players' expected true values $\{\hat{v}_i\}_{i \in \mathcal{N}}$ according to equation (4) and payoffs $\{\hat{u}_i\}_{i \in \mathcal{N}}$. Since equation (4) is identical to equation (2), we have $\hat{v}_i = v_i^\dagger$ for each player i . Thus, when the Simple Mechanism assigns objects to maximize its estimate of social surplus (i.e., to solve $\max_{\phi \in \Phi} \sum_{i \in \mathcal{N}} \hat{u}_i(\phi(i))$), it is actually maximizing social surplus (i.e., solving $\max_{\phi \in \Phi} \sum_{i \in \mathcal{N}} u_i(\phi(i)|\theta_i)$). Thus, the mixed assignment it produces is interim efficient.

We next wish to show that the Simple Mechanism is symmetric. Let i and k be two players who make the same report r and have the same type θ ; subscripts are suppressed for simplicity. Upon receiving report $r = (r_1, \dots, r_{i-1}, r, r_{i+1}, \dots, r_{k-1}, r, r_{k+1}, \dots, r_N)$, the Simple Mechanism generates estimates of players' expected true values $\{\hat{v}_l\}_{l \in \mathcal{N}}$ according to equation (4) and estimates of their payoffs $\{\hat{u}_l\}_{l \in \mathcal{N}}$. It also computes $\sigma = \arg \max_{\phi \in \Phi} \sum_{i \in \mathcal{N}} \hat{u}_i(\phi(i))$.

Partition σ as follows $\sigma_+ = \{\phi \in \sigma | u_i(\phi|\theta) > u_k(\phi|\theta)\}$, $\sigma_- = \{\phi \in \sigma | u_i(\phi|\theta) = u_k(\phi|\theta)\}$, and $\sigma_- = \{\phi \in \sigma | u_i(\phi|\theta) < u_k(\phi|\theta)\}$. Focus on the relationship between σ_+ and σ_- . For each $\phi \in \sigma_+$, we may construct an

³⁹We write $x \succ y$ to mean x is strictly preferred to y .

element $s(\phi) \in \sigma_-$ by having i and k swap their assignments, while leaving all others players' assignments unaltered, i.e., by setting $s(\phi)(i) = \phi(k)$, $s(\phi)(k) = \phi(i)$, and $s(\phi)(l) = \phi(l)$ for every other player l . This swap: (i) does not change the Simple Mechanism's estimate of the sum of players' payoffs since the mechanism estimates i and k have the same payoff (since they make the same report) and (ii) defines a bijection between σ_+ and σ_- .

Turning to i and k payoffs to the Simple Mechanism, we have $u_i(M_S(r)|\theta) =$

$$\begin{aligned} & \sum_{\phi \in \sigma_+} u_i(\phi|\theta) \frac{1}{|\sigma|} + \sum_{\phi \in \sigma_-} u_i(\phi|\theta) \frac{1}{|\sigma|} + \sum_{\phi \in \sigma_-} u_i(\phi|\theta) \frac{1}{|\sigma|} \\ = & \sum_{\phi \in \sigma_+} u_k(s(\phi)|\theta) \frac{1}{|\sigma|} + \sum_{\phi \in \sigma_-} u_k(\phi|\theta) \frac{1}{|\sigma|} + \sum_{\phi \in \sigma_-} u_k(s(\phi)|\theta) \frac{1}{|\sigma|} \\ = & \sum_{\phi \in \sigma_-} u_k(\phi|\theta) \frac{1}{|\sigma|} + \sum_{\phi \in \sigma_-} u_k(\phi|\theta) \frac{1}{|\sigma|} + \sum_{\phi \in \sigma_+} u_k(\phi|\theta) \frac{1}{|\sigma|} = u_k(M_S(r)|\theta). \end{aligned}$$

The second line follows from the fact that i and k swap assignments and therefore payoffs. The third line follows from reindexing: since s is bijective, we may replace each $\phi \in \sigma_+$ ($\phi \in \sigma_-$) with its corresponding element from σ_- (σ_+), implying $\sum_{\phi \in \sigma_+} u_k(s(\phi)|\theta) = \sum_{\phi \in \sigma_-} u_k(\phi|\theta)$ and $\sum_{\phi \in \sigma_-} u_k(s(\phi)|\theta) = \sum_{\phi' \in \sigma_+} u_k(\phi'|\theta)$. \square

We prove Proposition 3 in two steps. First, we establish an upper bound on any player's payoff to sending a report to the Simple Mechanism, when the other players tell the truth, as a function of her type and the number of players. Then, we show that this bound goes to her value of η as N gets large – Lemma A1. Second, we leverage this bound, the fact that a strategic liar never does worse than being unassigned, and a squeeze argument (similar to the one used to establish Observation 4) to show that each player's gain from strategic lying is small when N is large.

Lemma A1. A Uniform Upper Bound on $U_i^{M_S}(r_i|\theta_i)$.

Let \mathcal{V} be the product of finite grids and let $\alpha > 0$. Then, there is a function $g : \Theta \times \mathbb{N} \rightarrow \mathbb{R}$, such that, for each player i and each $\theta_i \in \Theta$:

(i) The function g bounds i 's payoff to reporting each $r_i \in \Theta$ from above, i.e., $U_i^{M_S}(r_i|\theta_i) \leq g(\theta_i, N)$.

(ii) The function g converges to i 's value of being unassigned as $N \rightarrow \infty$, i.e., $\lim_{N \rightarrow \infty} g(\theta_i, N) = u_i(\eta|\theta_i)$.

Proof. Consider player i . We proceed in four steps. First, we partition the type space of the other players in the right way. Second, on each of these partitions, we give upper bounds for i 's payoff given the behavior of the Simple Mechanism. Third, we average these upper bounds and write down a closed form expression of g . Fourth, we let $N \rightarrow \infty$ and show that g takes the claimed limits.

We begin by partitioning the type space Θ^{N-1} of the other players. We say a player is a **type E** if her type is $(v_1^\eta, v_{L_1}^1, v_{L_2}^2, \dots, v_{L_K}^K)$, i.e., if she (i) knows her true values, (ii) values being unassigned at the lowest possible value, and (iii) values every object at the highest possible value. (There are such players because \mathcal{V} is the product of finite grids, $\alpha > 0$, and $f_v > 0$.) Let E_l denote the event that there are l type E players besides i , i.e., $E_l = \{\theta_{-i} \in \Theta^{N-1} | \theta_{-i} \text{ specifies } l \text{ players have type E}\}$. The probability nature draws type E for a given player is $\alpha f_v(\bar{v}) > 0$, where $\bar{v} = (v_1^\eta, v_{L_1}^1, v_{L_2}^2, \dots, v_{L_K}^K)$. Thus, the probability of E_l is $\Pr(E_l) = \sum_{\theta_{-i} \in E_l} \Pr(\theta_{-i}) = \binom{N-1}{l} (\alpha f_v(\bar{v}))^l (1 - \alpha f_v(\bar{v}))^{N-1-l}$ due to independence. By the Binomial Theorem, $\sum_{l=0}^{N-1} \Pr(E_l) = 1$.

Next, we give upper bounds on $u_i(M_S(r_i, \theta_{-i}) | \theta_i)$ on each set E_l for each $r_i \in \Theta$ and each $\theta_i \in \Theta$. For l between 0 and K , it suffices to let 1 be this bound by the finite grid assumption. For $l > K$, we will establish that a bound is $u_i(\eta | \theta_i) + \frac{K}{l-K}$. It follows that

$$\begin{aligned}
U_i^M(r_i | \theta_i) &= \sum_{\theta_{-i} \in \Theta^{N-1}} u_i(M_S(r_i, \theta_{-i}) | \theta_i) \Pr(\theta_{-i}) \\
&= \sum_{l=0}^K \sum_{\theta_{-i} \in E_l} u_i(M_S(r_i, \theta_{-i}) | \theta_i) \Pr(\theta_{-i}) \\
&\quad + \sum_{l=K+1}^{N-1} \sum_{\theta_{-i} \in E_l} u_i(M_S(r_i, \theta_{-i}) | \theta_i) \Pr(\theta_{-i}) \\
&\leq \sum_{l=0}^K \sum_{\theta_{-i} \in E_l} 1 \Pr(\theta_{-i}) + \sum_{l=K+1}^{N-1} \sum_{\theta_{-i} \in E_l} (u_i(\eta | \theta_i) + \frac{K}{l-K}) \Pr(\theta_{-i}) \\
&= \sum_{l=0}^K \Pr(E_l) + \sum_{l=K+1}^{N-1} (u_i(\eta | \theta_i) + \frac{K}{l-K}) \Pr(E_l)
\end{aligned}$$

for each $r_i \in \Theta$ and each $\theta_i \in \Theta$. We take the last line to be our upper bound g ,

i.e., we set

$$g(\theta_i, N) = \sum_{l=0}^K \Pr(E_l) + \sum_{l=K+1}^{N-1} \left(u_i(\eta|\theta_i) + \frac{K}{l-K} \right) \Pr(E_l)$$

for each $\theta_i \in \Theta$. (The bound applies to every player, including player i , because players with the same type have the same payoff.)

We need to show that g takes the claimed limits. We begin by noting $\lim_{N \rightarrow \infty} \Pr(E_l) = 0$ for all $l < \infty$.⁴⁰ Thus, $\sum_{l=0}^K \Pr(E_l) \rightarrow 0$ as $N \rightarrow \infty$. To show that $\sum_{l=K+1}^{N-1} \left(u_i(\eta|\theta_i) + \frac{K}{l-K} \right) \Pr(E_l) \rightarrow u_i(\eta|\theta_i)$ for each $\theta_i \in \Theta$, we use a squeeze argument. Let $\delta > 0$, then there is a fixed N' such that (i) $\frac{K}{l-K} > \delta$ for all $l < N'$ and (ii) $\frac{K}{l-K} \leq \delta$ for all $l \geq N'$. Thus, for large N , we have

$$\begin{aligned} & \sum_{l=K+1}^{N-1} \left(u_i(\eta|\theta_i) + \frac{K}{l-K} \right) \Pr(E_l) \\ & \leq u_i(\eta|\theta_i) + \sum_{l=K+1}^{N'-1} \frac{K}{l-K} \Pr(E_l) + \delta \sum_{l=N'}^{N-1} \Pr(E_l) \\ & = u_i(\eta|\theta_i) + \underbrace{\sum_{l=K+1}^{N'-1} \frac{K}{l-K} \Pr(E_l) + \delta \left(1 - \sum_{l=0}^{N'-1} \Pr(E_l) \right)}_{h(N)}. \end{aligned}$$

Since $\lim_{N \rightarrow \infty} \Pr(E_l) = 0$, the finite sums in $h(N)$ go to zero as $N \rightarrow \infty$, implying $h(N) \rightarrow u_i(\eta|\theta_i) + \delta$. Thus, we have $\lim_{N \rightarrow \infty} g(\theta_i, N) \leq u_i(\eta|\theta_i) + \delta$ for

⁴⁰This follows from Sterling's Approximation and l'Hopital's rule. Let $p = \alpha f_v(\bar{v})$. For each finite l , Sterling's Approximation (see [Feller, 1968](#), p.52) lets us write

$$\lim_{N \rightarrow \infty} \Pr(E_l) = \lim_{N \rightarrow \infty} (N-1)^l (1-p)^{N-1} \left(1 + \frac{l}{N-1-l} \right)^{N-1-l} \sqrt{\frac{N-1}{N-1-l}} \frac{p^l}{l! e^l (1-p)^l}.$$

Since $\lim_{N \rightarrow \infty} \left(1 + \frac{l}{N-1-l} \right)^{N-1-l} = e^l$, $\lim_{N \rightarrow \infty} \sqrt{\frac{N-1}{N-1-l}} = 1$, and $\frac{p^l}{l! e^l (1-p)^l}$ is constant, the limit of $\Pr(E_l)$ is zero if and only if $\lim_{N \rightarrow \infty} (N-1)^l (1-p)^{N-1} = 0$. Let $\beta = 1/(1-p)$, so $(N-1)^l (1-p)^{N-1} = (N-1)^l / \beta^{N-1}$. Since $\beta > 1$ as $1-p \in (0, 1)$, we need to apply l'Hopital's rule. Doing so l times to gives

$$\lim_{N \rightarrow \infty} \frac{(N-1)^l}{\beta^N} = \lim_{N \rightarrow \infty} \frac{l!(N-1)^0}{\frac{(N-1)!}{(N-l-2)!} \beta^{N-1-l}}.$$

It is evident that the right-hand-side has limit zero, implying $\lim_{N \rightarrow \infty} (N-1)^l (1-p)^{N-1} = 0$.

all $\delta > 0$. Since $g(\theta_i, N) \geq u_i(\eta|\theta_i)$, it follows that $\lim_{N \rightarrow \infty} g(\theta_i, N) = u_i(\eta|\theta_i)$ for each $\theta_i \in \Theta$.

It only remains to show that, when $l > K$,

$$u_i(M_S(r_i, \theta_{-i})|\theta_i) \leq u_i(\eta|\theta_i) + \frac{K}{l-K}, \quad (6)$$

for each $r_i \in \Theta$ and each $\theta_{-i} \in E_l$. To do this, we need four facts.

Fact 0. For each $r_i \in \Theta$, for each $\theta_{-i} \in E_l$, and for each ϕ in the support of $M_S(r_i, \theta_{-i})$, there are at least $l - K$ unassigned type E players.

Proof. Simply there are K objects and $l > K$ type E players. Thus, at least $l - K$ of the type E players must be unassigned. \triangle

Fact 1. For each $r_i \in \Theta$ and each $\theta_{-i} \in E_l$, the Simple Mechanism assigns i an object $o \in \mathcal{O}$ with strictly positive probability only if it estimates that her value for said object is $v_{L_o}^o$. Thus, the Simple Mechanism assigns i to η if it estimates that her value for each object $o \in \mathcal{O}$ is strictly less than the maximum value for that object $v_{L_o}^o$.

Proof. Upon receiving report $r = (r_i, \theta_{-i})$, the Simple Mechanism generates estimates of players' expected true values $\{\hat{v}_l\}_{l \in \mathcal{N}}$ according to equation (4) and estimates of their payoffs $\{\hat{u}_l\}_{l \in \mathcal{N}}$. It also computes $\sigma = \arg \max_{\phi \in \Phi} \sum_{l \in \mathcal{N}} \hat{u}_l(\phi(l))$. Observe (i) that $\hat{u}_l(o) = u_l(o|\theta_l)$ for each player $l \neq i$ and each $o \in \mathcal{O}$ per truthful reporting and (ii) that $v_1^o \leq \hat{u}_l(o) \leq v_{L_o}^o$ for each player l and each $o \in \mathcal{O}$.⁴¹

We argue the fact by contradiction. Suppose i is assigned an object $o \in \mathcal{O}$ with strictly positive probability and the mechanism estimates that $\hat{u}_i(o) < v_{L_o}^o$. Then, there is a pure assignment ϕ in the support of $M_S(r_i, \theta_{-i})$ where i gets o . Since there is at least one unassigned type E under ϕ (by Fact 0), call her k , we can construct a new assignment ϕ' from ϕ by assigning o to k , by assigning η to i , and by leaving all other players' assignments alone – formally, $\phi'(k) = o$, $\phi'(i) = \eta$, and $\phi'(l) = \phi(l)$ for all $l \neq i, k$. We will show that $\sum_{l \in \mathcal{N}} \hat{u}_l(\phi'(l)) -$

⁴¹ Since equation (4) either returns an element of \mathcal{V} or averages over elements of \mathcal{V} and since every element of \mathcal{V} is bounded below by (v_1^η, \dots, v_1^K) and above by $(v_{L_\eta}^\eta, \dots, v_{L_K}^K)$ as it is the product of finite grids, we necessarily have $(v_1^\eta, \dots, v_1^K) \leq \hat{v}_l \leq (v_{L_\eta}^\eta, \dots, v_{L_K}^K)$ for each player l .

$\sum_{l \in \mathcal{N}} \hat{u}_l(\phi(l)) > 0$. Thus, ϕ does not maximize the Simple Mechanism's estimate of social welfare and therefore cannot be in the support of $M_S(r_i, \theta_{-i})$, a contradiction.

It remains to show that $\sum_{l \in \mathcal{N}} \hat{u}_l(\phi'(l)) - \sum_{l \in \mathcal{N}} \hat{u}_l(\phi(l)) > 0$. Since (i) $\hat{u}_i(o) < v_{L_o}^o$ and (ii) $\hat{u}_i(\eta) \geq v_1^\eta$, we have $v_{L_o}^o - \hat{u}_i(o) + \hat{u}_i(\eta) - v_1^\eta > 0$. Since k is a truthful type E , we have $\hat{u}_k(o) - \hat{u}_k(\eta) = v_{L_o}^o - v_1^\eta$. Hence,

$$\begin{aligned} \sum_{l \in \mathcal{N}} \hat{u}_l(\phi'(l)) - \sum_{l \in \mathcal{N}} \hat{u}_l(\phi(l)) &= \hat{u}_k(o) - \hat{u}_k(\eta) + \hat{u}_i(\eta) - \hat{u}_i(o) \\ &= v_{L_o}^o - \hat{u}_i(o) + \hat{u}_i(\eta) - v_1^\eta > 0. \quad \triangle \end{aligned}$$

Fact 2. For each $r_i \in \Theta$ and each $\theta_{-i} \in E_l$, the Simple Mechanism assigns η to i if it estimates that her value of η is strictly greater than v_1^η .

Proof. Upon receiving report $r = (r_i, \theta_{-i})$, the Simple Mechanism generates estimates of players' expected true values $\{\hat{v}_l\}_{l \in \mathcal{N}}$ according to equation (4) and estimates of their payoffs $\{\hat{u}_l\}_{l \in \mathcal{N}}$. It also computes $\sigma = \arg \max_{\phi \in \Phi} \sum_{l \in \mathcal{N}} \hat{u}_l(\phi(l))$. Recall (i) that $\hat{u}_l(o) = u_l(o|\theta_l)$ for each player $l \neq i$ and each $o \in \mathcal{O}'$ since these players report truthfully and (ii) that $v_1^o \leq \hat{u}_l(o) \leq v_{L_o}^o$ for each player l and each $o \in \mathcal{O}'$.

We argue this fact by contradiction using the same approach as in the proof of Fact 1. Suppose i is assigned an object $o \in \mathcal{O}$ with strictly positive and the mechanism estimates that $\hat{u}_i(\eta) > v_1^\eta$. Then, there is a pure assignment ϕ in the support of $M_S(r, \theta_{-i})$ where i gets o . Since there is at least one unassigned type E under ϕ , call her k , we can construct a new assignment ϕ' from ϕ by assigning o to k , by assigning η to i , and by leaving all other players' assignments alone – formally, $\phi'(k) = o$, $\phi'(i) = \eta$, and $\phi'(l) = \phi(l)$ for all $l \neq i, k$. We have

$$\begin{aligned} \sum_{l \in \mathcal{N}} \hat{u}_l(\phi'(l)) - \sum_{l \in \mathcal{N}} \hat{u}_l(\phi(l)) &= \hat{u}_k(o) - \hat{u}_k(\eta) + \hat{u}_i(\eta) - \hat{u}_i(o) \\ &= v_{L_o}^o - v_1^\eta + \hat{u}_i(\eta) - \hat{u}_i(o), \end{aligned}$$

because $\hat{u}_k(o) - \hat{u}_k(\eta) = v_{L_o}^o - v_1^\eta$. Since $\hat{u}_i(o) \leq v_{L_o}^o$ and $\hat{u}_i(\eta) > v_1^\eta$, we have $v_{L_o}^o - \hat{u}_i(o) + \hat{u}_i(\eta) - v_1^\eta > 0$. Thus, ϕ cannot be an element of $M_S(r_i, \theta_{-i})$, a contradiction. \triangle

Fact 3. For each $r_i \in \Theta$ and each $\theta_{-i} \in E_l$, if the Simple Mechanism assigns i an object $o \in \mathcal{O}$ with strictly positive probability, then the probability i gets o is bounded above by $\frac{1}{l-K}$. (Notice that this bound does *not* apply to η .)

Proof. Upon receiving report $r = (r_i, \theta_{-i})$, the Simple Mechanism generates estimates of players' expected true values $\{\hat{v}_l\}_{l \in \mathcal{N}}$ according to equation (4) and estimates of their payoffs $\{\hat{u}_l\}_{l \in \mathcal{N}}$. It also computes $\sigma = \arg \max_{\phi \in \Phi} \sum_{l \in \mathcal{N}} \hat{u}_l(\phi(l))$. Recall (i) that $\hat{u}_l(o) = u_l(o|\theta_l)$ for each player $l \neq i$ and each $o \in \mathcal{O}'$ and (ii) that $v_1^o \leq \hat{u}_l(o) \leq v_{L_o}^o$ for each player l and each $o \in \mathcal{O}'$.

Let $p_i(o)$ be the probability that i is assigned object $o \in \mathcal{O}$ in $M_S(r, \theta_{-i})$. Since the Simple Mechanism randomizes uniformly over σ , we have $p_i(o) = \sum_{\phi \in C_i(o)} \frac{1}{|\sigma|} = \frac{|C_i(o)|}{|\sigma|}$ where $C_i(o) = \{\phi \in \sigma | \phi(i) = o\}$ is the sub-set of σ where i gets o . (Since $p_i(o) > 0$ by hypothesis, $C_i(o)$ is non-empty.) Our goal is to show that $p_i(o) \leq \frac{1}{l-K}$, i.e., $|C_i(o)|(l-K) \leq |\sigma|$. We do this by showing that for each $\phi \in C_i(o)$, there are $l-K$ "unique" assignments in $\sigma \setminus C_i(o)$.

Pick a $\phi \in C_i(o)$ and let $S(\phi)$ be the set of unassigned type E players at ϕ . For each $k \in S(\phi)$ we can create an assignment ϕ' by swapping i and k 's assignments and leaving all other players' assignments alone, i.e., by setting $\phi'(k) = o$, $\phi'(i) = \eta$, and $\phi'(l) = \phi(l)$ for all $l \neq i, k$. Since there are at least $l-K$ players in $S(\phi)$ per Fact 0, we can create at least $l-K$ different assignments in this way. Let $P(\phi)$ be the set of assignments created in this fashion.

First, observe that $P(\phi) \subset \sigma \setminus C_i(o)$. It is trivial that $P(\phi) \subset \Phi \setminus C_i(o)$, since i is assigned η in every element of $P(\phi)$. So, we only need to show that $P(\phi) \subset \sigma$. This follows from Facts 1 and 2: since i receives object o with strictly positive probability in $M_S(r, \theta_{-i})$, it must be that $\hat{u}_i(\eta) = v_1^\eta$ and $\hat{u}_i(o) = v_{L_o}^o$. Thus, $\sum_{l \in \mathcal{N}} \hat{u}_l(\phi'(l)) - \sum_{l \in \mathcal{N}} \hat{u}_l(\phi(l)) = \hat{u}_{\phi'(o)}(o) - \hat{u}_{\phi'(o)}(\eta) + \hat{u}_i(\eta) - \hat{u}_i(o) = 0$ for any $\phi' \in P(\phi)$ since $\hat{u}_{\phi'(o)}(o) = v_{L_o}^o$ and $\hat{u}_{\phi'(o)}(\eta) = v_1^\eta$ as $\phi'(o)$ is truthful type E . It follows that $P(\phi) \subset \sigma$.

Second, observe that for two assignments ϕ and ϕ' in $C_i(o)$, we have that $P(\phi) \cap P(\phi') = \emptyset$. We argue this by contradiction. Suppose there is a $\tilde{\phi} \in P(\phi) \cap P(\phi')$. Since ϕ and ϕ' are different assignments, there is a player k for whom $\phi(k) \neq \phi'(k)$. (Notice $k \neq i$ as i has object o in both ϕ

and ϕ' .) There are three cases: $\tilde{\phi}(k) = o'$ for some $o' \in \mathcal{O} \setminus \{o\}$, $\tilde{\phi}(k) = \eta$, and $\tilde{\phi}(k) = o$. If $\tilde{\phi}(k) = o'$, then it must be that $\phi(k) = \phi'(k) = o'$ – since we do not chance the assignments of players who get objects other than o in the creation of $\tilde{\phi}$, the only way k could end up with o' is if she had it to begin with in both ϕ and ϕ' . If $\tilde{\phi}(k) = \eta$, then it must be that $\phi(k) = \phi'(k) = \eta$ since we never assign η to a player (except i) in the creation of $\tilde{\phi}$, unless said player has η to begin with. If $\tilde{\phi}(k) = o$, then it must be that $\phi(k) = \phi'(k) = \eta$ since we only give o to an unassigned type E player in the creation of $\tilde{\phi}$. In all three cases, $\phi(k) = \phi'(k)$, a contradiction.

Let $Q = \cup_{\phi \in C_i(o)} P(\phi)$. Since the $\{P(\phi)\}_{\phi \in C_i(o)}$ sets are disjoint by the second observation, $|Q| = \sum_{\phi \in C_i(o)} |P(\phi)|$. Since $|P(\phi)| \geq (l - K)$, we have $|Q| \geq (l - K)|C_i(o)|$. Since $Q \subset \sigma \setminus C_i(o)$ by the first observation, σ has at least $(l - K)|C_i(o)|$ elements. \triangle

This concludes the list of facts and the proofs thereof.

We now establish equation (6). Let $l > K$, let $\theta_{-i} \in E_l$, and, for each $r_i \in \Theta$, let $\hat{u}_i^{r_i}$ be the Simple Mechanism's estimate of i 's payoff function upon receiving the report (r_i, θ_{-i}) , see equation (4). We partition the set of all possible reports Θ into three parts Θ_1 , Θ_2 , and Θ_3 as follows: $\Theta_1 = \{r_i \in \Theta | \hat{u}_i^{r_i}(\eta) > v_1^\eta\}$, $\Theta_2 = \{r_i \in \Theta | \hat{u}_i^{r_i}(\eta) = v_1^\eta\} \cap \{r_i \in \Theta | \hat{u}_i^{r_i}(o) < v_{L_o}^o \text{ for all } o \in \mathcal{O}\}$, and $\Theta_3 = \{r_i \in \Theta | \hat{u}_i^{r_i}(\eta) = v_1^\eta\} \cap \{r_i \in \Theta | \hat{u}_i^{r_i}(o) = v_{L_o}^o \text{ for some } o \in \mathcal{O}\}$. In other words, Θ_1 is the set of reports that cause the mechanism to estimate that i 's value to η is strictly greater than v_1^η , Θ_2 is the set of reports that cause the mechanism to estimate that i 's value to η is v_1^η and her value to every object o is less than $v_{L_o}^o$, and Θ_3 is the set of reports cause the mechanism to estimate that i 's value to η is v_1^η and i 's value to some objects is maximal.

If $r_i \in \Theta_1 \cup \Theta_2$, then Facts 1 and 2 tells us the Simple Mechanism assigns η to i , so $u_i(M_S(r_i, \theta_{-i}) | \theta_i) = u_i(\eta | \theta_i)$. If $r_i \in \Theta_3$, then Fact 3 tells us that the probability i gets each object in \mathcal{O} is no more than $\frac{1}{l-K}$. Hence,

$$\begin{aligned} u_i(M_S(r_i, \theta_{-i}) | \theta_i) &= u_i(\eta | \theta_i) p_i(\eta) + \sum_{o \in \mathcal{O}} u_i(o | \theta_i) p_i(o) \\ &\leq u_i(\eta | \theta_i) + \sum_{o \in \mathcal{O}} u_i(o | \theta_i) \frac{1}{l-K} \leq u_i(\eta | \theta_i) + \frac{K}{l+1} \end{aligned}$$

for any $r_i \in \Theta_3$, where $p_i(o)$ is the probability i gets o in $M_S(r_i, \theta_{-i})$ and the last inequality is due to the finite grid assumption. Thus, across all reports

in $r_i \in \Theta$, we have $u_i(M_S(r_i, \theta_{-i})|\theta_i) \leq u_i(\eta|\theta_i) + \frac{K}{l-K}$. That is, equation (6) holds. \square

Proof of Proposition 3. Consider player i whose type is $\theta_i \in \Theta$. We begin by establishing that $u_i(\eta|\theta_i)$ is a lower bound on i 's earnings when she strategically lies or when she reports truthfully. Observe that $u_i(M_S(\theta_i, \theta_{-i})|\theta_i) \geq u_i(\eta|\theta_i)$ for all $\theta_{-i} \in \Theta^{N-1}$ by interim efficiency and individual rationality. It follows that $U_i^{M_S}(\theta_i|\theta_i) =$

$$\sum_{\theta_{-i} \in \Theta^{N-1}} u_i(M_S(\theta_i, \theta_{-i})|\theta_i) \Pr(\theta_{-i}) \geq \sum_{\theta_{-i} \in \Theta^{N-1}} u_i(\eta|\theta_i) \Pr(\theta_{-i}) = u_i(\eta|\theta_i).$$

Hence, $\max_{r_i \in \Theta} U_i^{M_S}(r_i|\theta_i) \geq U_i^{M_S}(\theta_i|\theta_i) \geq u_i(\eta|\theta_i)$ by optimality.

Now, let $\varepsilon > 0$. Then, by the second part of Lemma A1, there is a finite $N_{\theta_i} > 0$ such that $g(\theta_i, N) \leq u_i(\eta|\theta_i) + \varepsilon$ for all $N \geq N_{\theta_i}$. The first part of Lemma A1 and the prior paragraph then imply $u_i(\eta|\theta_i) \leq U_i^{M_S}(\theta_i|\theta_i) \leq u_i(\eta|\theta_i) + \varepsilon$ and $u_i(\eta|\theta_i) \leq \max_{r_i \in \Theta} U_i^{M_S}(r_i|\theta_i) \leq u_i(\eta|\theta_i) + \varepsilon$ for all $N \geq N_{\theta_i}$, i.e., i never gain more than ε by lying strategically when there are at least N_{θ_i} players. Thus, no player gains more than ε by strategically lying when there are at least $\bar{N} = \max_{\theta \in \Theta} N_{\theta}$ players (since players with the same type have the same payoff.) \square

Remark 1. It is clear from the Proofs of Lemma A1 and Proposition 3 that we can relax the assumption that \mathcal{V} is a product of finite grids to the assumption that, for each $o \in \mathcal{O}$, (i) there is a type of player who has the maximum value for o (among all players) and the minimum value for η and (ii) this player occurs with non-zero probability. We omit this abstract generalization because it adds little economic insight.

To extend the baseline model (of Section 2) to allow for object replication, we define the **replication extension** using the baseline model's primitives – \mathcal{O} , \mathcal{V} , f_v , α , and N – as well as the linking function $\psi(N)$. This extension makes two key modifications to the baseline model. First, the set of objects is $\cup_{l=1}^{\psi(N)} \mathcal{O}$ and a pure assignment ϕ is a map from \mathcal{N} into $\{\eta\} \cup_{l=1}^{\psi(N)} \mathcal{O} = \mathcal{O}'_{\psi(N)}$ subject to the usual restrictions. We write

$$\mathcal{O}'_{\psi(N)} = \{\eta, 1_1, 1_2, \dots, 1_{\psi(N)}, 2_1, \dots, 2_{\psi(N)}, \dots, K_1, \dots, K_{\psi(N)}\},$$

where o_l refers to the l -th copy of object o . Second, copies of the objects are indistinguishable. Thus, nature draws true values and types as in the baseline

version (of the model) and players' expected true values and payoffs are formed as in the baseline version – i.e., using equation (2). The true values, expected true values, and payoffs are then extended from \mathcal{O}' to $\check{\mathcal{O}}'_{\psi(N)}$ by treating the copies of each object identically. Thus, when $v_i^\dagger = (v_{\eta i}^\dagger, v_{1i}^\dagger, \dots, v_{K_i}^\dagger)$ is player i 's expected true values after learning her type θ_i , the copies of object o , i.e., $o_1, \dots, o_{\psi(N)}$, each have an expected true value of v_{oi}^\dagger . Consequently, (i) player i 's payoff is $u_i(o_l | \theta_i) = v_{oi}^\dagger$ for each object $o \in \mathcal{O}$ and each copy $l \in \{1, \dots, \psi(N)\}$ and (ii) her payoff to η is $u_i(\eta | \theta_i) = v_{\eta i}^\dagger$. Given these payoffs, concepts like interim efficiency and ε -Bayesian incentive compatibility extend naturally.

We also extend the Simple Mechanism to allow for object replication by making it aware of these payoff extensions. Specifically, for each player i , after employing equation (4) and estimating \hat{u}_i on \mathcal{O}' using i 's report $r_i \in \Theta$, it extends \hat{u}_i to $\check{\mathcal{O}}'_{\psi(N)}$ by setting $\hat{u}_i(o_l | \theta_i) = \hat{u}_i(o | \theta_i)$ for each $o \in \mathcal{O}$ and each copy $l \in \{1, \dots, \psi(N)\}$. It then computes the set of pure assignments that maximize the sum of the estimated payoffs and randomizes in the usual fashion.

Consider a version of the baseline model, which we call the **T -replication case**, where T denotes the number of replications. The case is defined via the primitives

$$\check{\mathcal{O}}_T = \cup_{l=1}^T \mathcal{O} = \{1_1, 1_2, \dots, 1_T, 2_1, \dots, 2_T, \dots, K_1, \dots, K_T\},$$

$\check{\mathcal{V}} = \mathcal{V}_\eta \times_{o \in \mathcal{O}} \check{\mathcal{V}}_o(T)$ where $\check{\mathcal{V}}_o(T) = \{(x_1, \dots, x_T) \in \mathbb{R}^T | \exists v \in \mathcal{V}_o \text{ s.t. } x_1 = x_2 = \dots = v\}$ for each $o \in \mathcal{O}$ (per the finite grid assumption),

$$f_{\check{v}}(v_\eta, v_{1_1}, \dots, v_{2_1}, \dots, v_{K_1}, \dots, v_{K_T}) = f_v(v_\eta, v_{1_1}, v_{2_1}, \dots, v_{K_1})$$

for each $(v_\eta, v_{1_1}, \dots, v_{2_1}, \dots, v_{K_1}, \dots, v_{K_T}) \in \check{\mathcal{V}}$, α , and N . In this case, players' true values are defined directly on $\check{\mathcal{O}}'_T = \{\eta\} \cup \check{\mathcal{O}}_T$ (by treating copies of objects identically), rather than on \mathcal{O} and then extended.⁴² It is clear that players may have the same true values in the T -replication case and the replication extension when $T = \psi(N)$.

When $T = \psi(N)$, there is a natural bijection between the type-spaces of these environments: (i) a player's type is $(v_\eta, v_1, v_2, \dots, v_K)$ in the replication

⁴² We give the replication extension to avoid the technical challenges of modifying players types to account for growth in the number of copies in the T -replication case.

extension if and only if her type in the T -replication case is

$$(\eta, \underbrace{v_1, \dots, v_1}_{T \text{ copies}}, \underbrace{v_2, \dots, v_2}_{T \text{ copies}}, \dots, \underbrace{v_K, \dots, v_K}_{T \text{ copies}})$$

and (ii) a player's type is \preceq in the replication extension if and only if her type in the T -replication case is \preceq' , where \preceq and \preceq' are such that, for any $o, o' \in \mathcal{O}'$ and any $l, l' \in \{1, \dots, T\}$, we have (i) $o \preceq o'$ if and only if $o_l \preceq' o'_{l'}$ and (ii) $o \not\preceq o'$ if and only if $o_l \not\preceq' o'_{l'}$. Let $h : \Theta \rightarrow \tilde{\Theta}$ denote this bijection, where Θ denotes the replication extension type-space and $\tilde{\Theta}$ denotes the T -replication case type-space. It is clear that player i with type θ_i has the same expected true values and payoffs in the replication extension as she does in the T -replication case when her type is $h(\theta_i)$ and vice versa. (Also, let $h(\theta_1, \theta_2, \dots, \theta_N) = (h(\theta_1), h(\theta_2), \dots, h(\theta_N))$ and $h(r_1, r_2, \dots, r_N) = (h(r_1), \dots, h(r_N))$ for any $(\theta_1, \dots, \theta_N)$ or (r_1, \dots, r_N) in Θ^N .)

The payoff equivalence implies that if player i , with type θ_i , makes report r_i to the Simple Mechanism in the replication extension, then she has the same payoff to the mechanism as in the T -replication case when she has type $h(\theta_i)$ and makes report $h(r_i)$ to the mechanism. In symbols, $U_i^{Ms}(r_i|\theta_i) = \tilde{U}_i^{Ms}(h(r_i)|h(\theta_i))$ for all $r, \theta_i \in \Theta$ and each player i when $T = \psi(N)$, where $U_i^{Ms}(r_i|\theta_i)$ denotes player i 's payoff, given her type $\theta_i \in \Theta$, to making report $r_i \in \Theta$ to the Simple Mechanism in the replication extension and (in an abuse of notation) $\tilde{U}_i^{Ms}(r_i|\theta_i)$ denotes i 's payoff, given her type $\theta_i' \in \tilde{\Theta}$, to making report $r_i' \in \tilde{\Theta}$ to the Simple Mechanism in the T -replication case.

To better understand this equivalence, suppose an arbitrary player k reports r_k to the Simple Mechanism in the replication extension and $h(r_k)$ in the T -replication case. If these reports were her types, it is evident that her expected true values would be the same in both environments. Thus, the Simple Mechanism estimates she has the same payoff in both environments. It follows that, for any joint vector of players' reports r in the replication extension, the mechanism has the same objective function and thus randomizes over the same pure assignments as in the T -replication case when players report $h(r)$. Given this, given the payoff equivalence, and given the fact that each joint vector of types θ in the replication extension has the same probability as the joint type vector $h(\theta)$ in the T -replication case (and vice versa), it is straightforward to show that $U_i^{Ms}(r_i|\theta_i) = \tilde{U}_i^{Ms}(h(r_i)|h(\theta_i))$ for each player i

and each $r_i, \theta_i \in \Theta$.⁴³

Proof of Proposition 4. To establish the first part of the proposition, it suffices to provide an example where every interim efficient and symmetric mechanism is not ε -Bayesian incentive compatible. We do this in Example A (below), while further elaborating on the intuition for the failure of ε -Bayesian incentive compatibility.

Fix $\varepsilon > 0$ and consider player i . To establish the second part of the proposition, it suffices to show that there are (i) a linking function $\bar{\psi}$ and (ii) a \bar{N} such that $\max_{r_i \in \Theta} U_i^{Ms}(r_i|\theta_i) \leq U_i^{Ms}(\theta_i|\theta_i) + \varepsilon$ for all $N \geq \bar{N}$ for each $\theta_i \in \Theta$ in the replication extension. (This is sufficient because all players with the same type have the same payoff.)

To do this, first consider the T -replication case. Since this case is a version of the baseline model, Proposition 3 holds and gives there is some N_T such that $\max_{r_i \in \Theta} \dot{U}_i^{Ms}(r_i|\theta_i) \leq \dot{U}_i^{Ms}(\theta_i|\theta_i) + \varepsilon$ for all $N \geq N_T$ for each $\theta_i \in \Theta$. Thus, there is a sequence of thresholds $N_1, N_2, \dots, N_T, \dots$ that ensure ε -Bayesian incentive compatibility in the 1, 2, \dots , T , \dots replication cases respectively. It is without loss to suppose $N_1 < N_2 < \dots < N_T < \dots$. We use this sequence of thresholds to construct $\bar{\psi}$ as follows,

$$\bar{\psi}(n) = \begin{cases} 1 & \text{if } n < N_2 \\ 2 & \text{if } N_2 < n \leq N_3 \\ \vdots & \vdots \\ T & \text{if } N_T < n \leq N_{T+1} \\ \vdots & \vdots \end{cases}.$$

It is evident that this function is weakly monotone on \mathbb{N} and strictly monotone on an infinite sub-set of \mathbb{N} , implying $\lim_{n \rightarrow \infty} \bar{\psi}(n) = \infty$.

Returning to the replication extension, when the number of copies of each object is determined by $\bar{\psi}(N)$, we claim that $\bar{N} = N_1$ is the required threshold. To see this, pick some number of players $N \geq \bar{N}$ and observe that $N \in [N_{T'}, N_{T'+1}]$ for some $T' \geq 1$. Since (i) $\max_{r_i \in \Theta} \dot{U}_i^{Ms}(r_i|\theta_i) \leq \dot{U}_i^{Ms}(\theta_i|\theta_i) + \varepsilon$ for each $\theta_i \in \Theta$ in the T' -replication case and since (ii)

$$U_i^{Ms}(r_i|\theta_i) = U_i^{Ms}(h(r_i)|h(\theta_i))$$

⁴³ We formally establish this equality in the Online Appendix.

for all $r_i, \theta_i \in \Theta$ because $\bar{\psi}(N) = T'$, we have, $\max_{r_i \in \Theta} U_i^{Ms}(r_i | \theta_i) =$

$$\max_{r_i \in \Theta} U_i^{Ms}(h(r_i) | h(\theta_i)) \leq \ddot{U}_i^{Ms}(h(\theta_i) | h(\theta_i)) + \varepsilon = U_i^{Ms}(\theta_i | \theta_i) + \varepsilon,$$

for each $\theta_i \in \Theta$ in the replication extension. \square

Example A. Failure of ε -Bayesian Incentive Compatibility with an Arbitrary Linking Function.

Suppose there are $N \geq 2$ players and there is a single object o with $\psi(N) = \lfloor N/2 \rfloor$ copies. Further, suppose $\mathcal{V} = \{0\} \times \{\frac{1}{2}, 1\}$, with $f_v(\cdot) = \frac{1}{2}$ and $\alpha = 1$. That is, players observe their values and there are two equally likely types of players: type “As” who have values $(v_\eta, v_o) = (0, 1)$ and type Bs who have values $(v_\eta, v_o) = (0, \frac{1}{2})$.

After receiving reports, it is evident that any interim efficient mechanism assigns the object to type A reporters before assigning any remaining copies to type B reporters. Furthermore, it is readily verified that any symmetric mechanism endows each type A reporter with the same probability of receiving a copy of the object, likewise for each type B reporter. Let M denote an interim efficient and symmetric mechanism.

To show that there is no symmetric, interim efficient, and ε -Bayesian incentive compatible mechanism in this example, focus on the truth-telling incentives player for 1 when she is a type B. Let A_l be the event there are $l \in \{0, 1, \dots, N-1\}$ other players who are type A. If all other players report truthfully, then given l the probability a type A gets an object is $\frac{\min\{\lfloor N/2 \rfloor, l\}}{l}$ and the probability a type B gets an object is $\frac{\max\{\lfloor N/2 \rfloor - l, 0\}}{N-l}$.⁴⁴ Thus, player 1’s payoff to telling the truth to M is

$$U_1^M(B|B) = \frac{1}{2} \sum_{l=0}^{N-1} \frac{\max\{\lfloor N/2 \rfloor - l, 0\}}{N-l} \binom{N-1}{l} \left(\frac{1}{2}\right)^{N-1},$$

where $\binom{N-1}{l} \left(\frac{1}{2}\right)^{N-1}$ is the probability of l type A players and $N-l$ type B players. Alternatively, if player 1 were to lie and report she is a type A, she

⁴⁴ When there are l type As and $\lfloor N/2 \rfloor$ objects, either (i) every type A gets an object with certainty or (ii) they enter into a uniform lottery for these objects (by symmetry). After the type As are assigned, there are $\max\{\lfloor N/2 \rfloor - l, 0\}$ objects for the type Bs, so either (i’) every type B get an object with certainty or (ii’) the type Bs enter a uniform lottery for the remaining objects. Thus, $\frac{\min\{\lfloor N/2 \rfloor, l\}}{l}$ is the probability a type A is assigned the object and $\frac{\max\{\lfloor N/2 \rfloor - l, 0\}}{N-l}$ is the probability a type B is assigned an object (since player 1 is a type B).

would get the object with probability $\frac{\min\{\lfloor N/2 \rfloor, l+1\}}{l+1}$ given l . Thus, her payoff to lying to M is

$$U_1^M(A|B) = \frac{1}{2} \sum_{l=0}^{N-1} \frac{\min\{\lfloor N/2 \rfloor, l+1\}}{l+1} \binom{N-1}{l} \left(\frac{1}{2}\right)^{N-1}.$$

Hence, player 1's gain to lying is

$$\begin{aligned} U_1^M(A|B) - U_1^M(B|B) &= \\ \frac{1}{2} \sum_{l=0}^{N-1} \left(\frac{\min\{\lfloor N/2 \rfloor, l+1\}}{l+1} - \frac{\max\{\lfloor N/2 \rfloor - l, 0\}}{N-l} \right) \binom{N-1}{l} \left(\frac{1}{2}\right)^{N-1} &= \\ \frac{1}{2} \sum_{l=0}^{\lfloor N/2 \rfloor - 1} \left(1 - \frac{\lfloor N/2 \rfloor - l}{N-l} \right) \binom{N-1}{l} \left(\frac{1}{2}\right)^{N-1} + & \\ \frac{1}{2} \sum_{l=\lfloor N/2 \rfloor}^{N-1} \left(\frac{\lfloor N/2 \rfloor}{l+1} \right) \binom{N-1}{l} \left(\frac{1}{2}\right)^{N-1}. & \end{aligned}$$

It is clear that player 1 always gains from lying; this is illustrated in Figure A.1, which plots the value of this difference for N from 2 to 500. In fact, algebra shows (i) that the difference is always at least $\frac{1}{4}$ and (ii) that $\lim_{N \rightarrow \infty} U_1^M(A|B) - U_1^M(B|B) = \frac{1}{2}$. In other words, not only does player 1 gain a strictly positive, minimum amount from lying, but her gain actually increases with the number of players! It follows that, any symmetric and interim efficient mechanism M is not ε -Bayesian incentive compatible. The intuition for this outcome is that, with high probability, there are approximately as many objects as there are type A players. Thus, any symmetric and interim efficient mechanism typically assigns (almost) all of the objects to the type A s, leaving few to none for the type B s. Thus, a type B can obtain the object with high probability by lying. All of this is possible because replication negates competition for the object among the type A s. If competition were present among the type A s, then a type B 's gains to lying would evaporate and ε -Bayesian incentive compatibility would obtain (since no type A has incentive to lie). \triangle

Example B. ε -Bayesian Incentive Compatibility when Replication Occurs Slowly.




Figure 1: Example A, $U_1^M(A|B) - U_1^M(B|B)$ for $N \in \{2, 3, \dots, 500\}$

Consider the same setup as Example A, but instead suppose that $\psi(N) = \lfloor \ln(N)/2 \rfloor$ and $N \geq 8$.⁴⁵ Then, algebra shows that player 1's gain to lying to the Simple Mechanism is

$$U_1^{Ms}(A|B) - U_1^{Ms}(B|B) = \frac{1}{2} \sum_{l=0}^{N-1} \left(\frac{\min\{\lfloor \ln(N)/2 \rfloor, l+1\}}{l+1} - \frac{\max\{\lfloor \ln(N)/2 \rfloor - l, 0\}}{N-l} \right) \binom{N-1}{l} \left(\frac{1}{2}\right)^{N-1}.$$

While this difference is initially positive, it is evident that it shrinks to zero as $N \rightarrow \infty$. This behavior is illustrated in Figure A.2 which plots the gain to lying for N from 8 to 500. The intuition is that, unlike in Example A, the number of objects increases sufficiently slowly so that competition intensifies among the type As, limiting player 1's and thus any type B's gain to misrepresentation. Since no type A player has incentive to lie, ε -Bayesian incentive compatibility obtains. \triangle

References

Abdulkadiroglu, A., Che, Y., & Yasuda, Y. (2008). Expanding "Choice" in School Choice. Working Paper, Duke University.

⁴⁵ We take $N \geq 8$ so that $\psi(N) \geq \lfloor \ln(8)/2 \rfloor = 1$.




Figure 2: Example B, $U_1^{Ms}(A|B) - U_1^{Ms}(B|B)$ for $N \in \{8, 9, \dots, 500\}$

- Abdulkadiroglu, A., Che, Y., & Yasuda, Y. (2011, February). Resolving Conflicting Preferences in School Choice: The “Boston Mechanism” Reconsidered. The American Economic Review, 101(1), 399–410.
- Abdulkadiroglu, A., & Sonmez, T. (1998). Random Serial Dictatorship and the Core From Random Endowments in House Allocation Problems. Econometrica: Journal of the Econometric Society, 66(3), 689–701.
- Abdulkadiroglu, A., & Sonmez, T. (2003). School Choice: A Mechanism Design Approach. American Economic Review, 93(3), 729–747.
- Azevedo, E., & Budish, E. (2018). Strategy proofness in the large. Working Paper, University of Pennsylvania.
- Bogomolnaia, A., & Moulin, H. (2001). A New Solution to the Random Assignment Problem. Journal of Economic Theory, 100(2), 295–328.
- Bu, N. (2014). Characterizations of the Sequential Priority Rules in the Assignment of Object Types. Social Choice and Welfare, 43(635–645).
- Budish, E. (2011, December). The Combinatorial Assignment Problem: Approximate Competitive Equilibrium from Equal Incomes. Journal of Political Economy, 119(6), 1061–1103.
- Budish, E., Che, Y., Kojima, F., & Milgrom, P. (2013). Designing Random Allocation Mechanisms: Theory and Applications. American Economic Review, 103(2), 585–623.
- Che, Y., & Tercieux, O. (2013). Efficiency and Stability in Large Matching Markets. Working Paper, Columbia University.
- Ehlers, L., Majumdar, D., Mishra, D., & Sen, A. (2014). Continuous Cardinal

- Incentive Compatible Mechanisms are Ordinal. Working Paper, University of Montreal.
- Feller, W. (1968). An Introduction to Probability Theory and its Applications. John Wiley and Sons.
- Gale, D., & Shapley, L. (1962). College Admissions and the Stability of Marriage. The American Mathematical Monthly, 69(1), 9–15.
- Hafalir, I., & Miralles, A. (2014). Welfare Maximizing Assignment of Agents to Hierarchical Positions. Working Paper, Carnegie Mellon University.
- He, Y., Miralles, A., Pycia, M., & Yan, J. (2018). A Pseudo-Market Approach to Allocation with Priorities. American Economic Journal: Microeconomics, 10(3), 272-314.
- Hylland, A., & Zeckhauser, R. (1979). The Efficient Allocation of Individuals to Positions. Journal of Political Economy, 87(2), 293-314.
- Kasajima, Y. (2013). Probabilistic Assignment of Indivisible Goods with Single-Peaked Preferences. Social Choice and Welfare, 41, 203-215.
- Liu, Q., & Pycia, M. (2016). Ordinal Efficiency, Fairness, and Incentives in Large Markets. Working Paper, UC Los Angeles.
- Miralles, A. (2008). School Choice: The Case for the Boston Mechanism. Working Paper, Autonomous University of Barcelona.
- Miralles, A., & Pycia, M. (2014). Cardinal Efficiency in Assignments Without Transfers. Working Paper, UC Los Angeles.
- Nguyen, T., Peivandi, A., & Vohra, R. (2014). One-Sided Matching with Limited Complementarities. Working Paper, University of Pennsylvania.
- Nicolo, A., & Rodriguez-Alvarez, C. (2013). Incentive Compatibility and Feasibility Constraints in Housing Markets. Social Choice and Welfare, 41, 625-635.
- Pycia, M. (2011, June). The Cost of Ordinality. Working Paper, UC Los Angeles, 1–4.
- Roth, A., Sonmez, T., & U., U. (2004). Kidney Exchange. Quarterly Journal of Economics, 119(2), 457-488.
- Roth, A., & Sotomayor, M. (1990). Two-sided matching: A study in game-theoretic modeling and analysis. Cambridge University Press.
- Shapley, L., & Scarf, H. (1974). On Cores and Indivisibility. Journal of Mathematical Economics, 1(1), 23–37.
- Sonmez, T., & Unver, U. (2010). House Allocation with Existing Tenants: A Characterization. Games and Economic Behavior, 69, 425-445.
- Zhou, L. (1990). On a Conjecture by Gale About One-Sided Matching Problems. Journal of Economic Theory, 52(1), 123–135.