Information Design in Matching Markets

Sulagna Dasgupta

March 2, 2022

Abstract

In the (one-sided) matching problem, objects are allocated to agents without monetary transfers, based on agents' preferences. However, agents may not always know, a priori, their preferences over the objects, because they do not have enough information. In this context, I try to answer the question: How should a benevolent planner optimally reveal information to the agents to maximize welfare, in an environment where agents have no private information to start with? As a benchmark, I first show that when using any of the standard strategyproof ordinal mechanisms, such as Deferred Acceptance, Serial Dictatorship, Random Priority or Top Trading Cycle, letting each agent know his true ordinal ranking over the objects is almost never a social welfare-maximizing information policy. By way of a partial solution, I then propose a simple signal I call the Object Recommendation (OR) Signal. Under independent agent priors satisfying a mild regularity condition, I show that, when agents' a priori relative preferences over the objects are "not too strong", the OR Signal, used together with any of the aforementioned standard mechanisms, not only maximizes welfare, but achieves first-best, i.e. the unconstrained maximum total ex-ante welfare.

1 Introduction

1.1 Motivation

Consider the problem of allocating a set of objects among a set of agents without monetary transfers, commonly known as the *matching/object allocation* problem. Examples of such a setting include school choice,college admissions, house allocation, allocation of tasks across employees etc. In such settings it is common for participants to have imperfect information about their own preferences over the choices on offer. While mechanism design techniques have been extensively applied to improve welfare in these contexts, traditional models usually assume that agents have perfect information about their own preferences. Hence the usual notions of efficiency and welfare maximization commonly used in the matching literature fail in environments where participants are not sure about their own preferences. Recognizing this issue of imperfect information in the context of school choice, Glazerman et al. (2018) note that "If confusion [regarding the information on schools] results in uninformed decision-making then parents might not select schools that best fit their children's needs and the possible benefits of school choice could be undermined." Gallup (2017)'s 2017 survey found if they had to do it over again, the majority of Americans (51 %) who pursued a postsecondary education would change their degree type, institution or major - indicating the lack of the effective information available at the time of admissions to colleges.

In this paper, we analyze the standard object allocation problem - but assume that the agents do not know their true preferences over objects. A benevolent planner designs a test to optimally reveal information to agents about their preferences, in order to maximize social welfare, while also ensuring truth-telling by them after they have taken the revealed information into account.

It is perhaps intuitive to conjecture, that any loss of welfare arising out of the particularities of the informational setting must be mainly due to agents not knowing their own preferences well enough. For example, Chen and He (2017) - who demonstrate that students typically overpay for information in school choice in a lab setting - echo this intuition by suggesting that education authorities can improve welfare by "providing more information" to students, because this would save "socially wasteful costs of information acquisition". In this context, I first fully characterize the conditions under which the (arguably) most intuitive information policy - giving each agent full information about his preference - is optimal (Proposition 1). I then use this characterization to show that for any (non-constant) ordinal mechanism, full information provision is suboptimal for any generic prior distribution, and also, if the mechanism is strategyproof (e.g. popular mechanisms like Deferred Acceptance, Serial Dictatorship, Random Priority, Top Trading Cycles etc.) and the prior has full support it is never optimal to reveal full information about their own preferences to the agents (Theorem 1).

Several possibilities arise when the aforementioned full support assumption is relaxed, which gives rise to optimality of the full information signal in various specialized domains. For example, when there are just two objects (with potentially multiple copies) and the setting is *dichotomous* - i.e. objects can only be *acceptable* or *unacceptable* to agents (Bogomolnaia and Moulin, 2004; Bogomolnaia et al., 2005) - full information is a welfare-maximizing signal to use, as long as the (ordinal) mechanism used is strategyproof and efficient (Proposition 7). The ordinal rank-value preference domain - a cardinal domain of preferences where each ordinal rank is associated with a unique value regardless of the object occupying that rank in any agent's ordinal preference ranking (Featherstone, 2011, 2014) - offers yet another avenue for possibilities with the full information signal. Turns out, in any ordinal rank-value domain, under an independent and uniform prior, any pointwise welfare maximizing mechanism is incentive compatible, and therefore the full information signal is optimal (Proposition 8).

As a natural next step suggested by the main negative result, I then investigate the optimal information policy, for a special case - when First Best, i.e. the pointwise maximum aggregate welfare, is achievable. I introduce a simple and intuitive signal, which consists of simply recommending each agent to pick the object the planner would most like him to pick under the pointwise aggregate welfare maximizing allocation. I call this signal the *Object Recommendation (OR)* signal. I show that the OR signal can be seen as the canonical signal characterizing implementability of the First Best by the Agent-Proposing Deferred Acceptance mechanism when objects have no priorities over agents, which is also the same as the Random Serial Dictatorship mechanism (Theorem 2).

In the cases considered in this paper, the OR signal works through a simple process - by making each agent's recommended object his most-preferred object, *a posteriori*. Naturally, it is not possible to implement this for all priors. Priors for which this is possible are said to satisfy the *posterior top* property. I show that when agents are symmetric, a four-way equivalence holds between the posterior top property and implementability of First Best by Serial Dictatorship, Random Serial Dictatorship and *any* ordinal mechanism satisfying a property I call "weak efficiency" - the mild efficiency criterion which requires that whenever it is feasible to allocate every agent his most-preferred object, it is allocated (Proposition 2).

However, the posterior top property is an endogenously defined property of the prior. Hence, subsequently I go on to provide conditions on the primitives which deliver implementability of the First Best by the large class of weakly efficient mechanisms. First, I consider the case of two objects, to build intuition, which I later generalize to the case of any finite set of objects. An additional canonical feature of the OR signal becomes salient in case of two objects. I show that First Best is implementable by any weakly efficient mechanism if and only if the prior is such that, the OR signal can *reverse* each agent's prior preference, whenever it is necessary to do so for aggregate welfare maximization (Proposition 3). Interestingly, if agents are, in addition, symmetric, the OR signal is equivalent to telling the agents their *rank* in the realized distribution of the agents' relative preferences over the two objects (Lemma 3), and always weakly improves welfare over not providing any information to agents (Proposition 4). Taken together, these two insights tell us that in the symmetric case, telling agents only their ranks in the relative-preference distribution is always better than not revealing any information and can sometimes (depending on the prior) achieve the highest possible aggregate welfare. Hence, these findings provide a welfare-based possible explanation of the ubiquity of rank-based information policies observed in practice.

Finally, I generalize the above insights from the two-object case to the case of any finite set of objects, and provide sufficient conditions on the priors for first best to be implementable by any weakly efficient ordinal mechanism. In particular, under a mild regularity condition on the prior, I show that when agents have independent preferences and do not have "too strong" a priori preferences over the objects, the OR signal implements first best for any weakly efficient mechanism (Theorem 3).

Several extensions are considered, in which I show that the main results presented in this paper are robust to different specifications of the planner's objective function (Proposition 5), restricting the set of priors to independent priors (Theorem 1' in Section 5.2) and restricting the set of possible allocations (Corollary 5) etc. The implications of jointly designing the mechanism and information structure are also considered (Section 5.5). I conclude by discussing avenues for future research.

1.2 Related literature

To the best of my knowledge, the current paper is the first to consider information design as a tool to improve welfare in the context of the matching problem. It contributes to the vast theoretical literature on matching in which it is typically assumed that agents know their preferences perfectly (Gale and Shapley (1962), Roth and Sotomayor (1992), Abdulkadiroğlu and Sönmez (2003), Bogomolnaia and Moulin (2001)). A relatively small part of this theoretical literature does investigate the role of incomplete information in matching settings. The incompleteness of information about one's own preferences in these models arise mainly from interdependence of values, either across the two sides of the market or within (Bikhchandani (2014), Lazarova and Dimitrov (2017), Liu et al. (2014), Chakraborty et al. (2010)). In recent times the school choice and college admissions literature has started acknowledging the role of imperfect private information in these settings even without interdependence. Examples of works in this category include Corcoran et al. (2018), Grenet et al. (2019), Bade (2015), Noda (2021), Chen and He (2017), and Immorlica et al. (2020). However, the focus of this literature has been mainly on the case where information acquisition is a strategic choice by *agents* and therefore on types of information flows incentivized by popular school choice/college admission mechanisms, rather than on strategic information provision by a benevolent designer.

While the emphasis of the above literature has been on providing agents *more* information, some of the applied literature does find what is theoretically expected - that in matching settings, capacity constraints may diminish the potential positive impact of providing agents with more information at both an individual (Neilson et al., 2019) and societal level (Sagaceta, 2020). Using a field experiment and simulations, Neilson et al. (2019) show that, while more information shifts parents' choices towards better schools - capacity constraints play an important role in reducing the impact of this informational intervention. Sagaceta (2020) studies the effects of implementing an AI-based recommendation algorithm within the Chilean school choice system. The goal of the algorithm is to recommend to each family the school that is predicted to be the best fit for their child, based on a range of factors. Unsurprisingly, the author finds, based on simulations, that this intervention has a negative impact on the utilitarian social welfare, which, the author explains, is the result of increased competition created by more informed agents.

Related to the above, our paper also speaks to the literature on “smart” recommender systems in the context of public services. Recommender systems are algorithms which use the observable attributes of items and/or users to generate personalized recommendations (Almazro et al., 2010).[1] Recommender systems are ubiquitious in e-commerce applications such as Netflix and Amazon (Gomez-Uribe and Hunt, 2015; Linden et al., 2003). But studies of their potential use in the context of public services - especially those involving economic applications - have been limited in number (Cortés-Cediel et al., 2017). Sagaceta (2020) is a notable exception. As described, most of this literature - including Sagaceta (2020) - focuses on improving the quality of recommendations at an individual level. Our paper provides an economic theory-based foundation for designing such recommendation systems with a welfarist goal.

Lastly, this paper contributes to the literature on applications of information design in games. The literature on information design - arguably launched by Kamenica and Gentzkow (2011)’s influential paper - studies the strategic revelation of information to agents by a designer, to further her own objectives. Bergemann and Morris (2016) takes Kamenica and Gentzkow (2011)’s framework for the single-agent problem to games, and formulates a Myersonian approach to Bayes Nash information design. This approach is based on a notion of correlated equilibrium under incomplete information - called *Bayes correlated equilibrium (BCE)* - which characterizes all possible Bayes Nash equilibrium outcomes that could arise under all information structures. Notable applications of this approach include Taneva (2019), who uses the above approach to study symmetric games in a binary setting. The current paper adds to the extensive literature on applications of information design in strategic settings which has subsequently developed, e.g. in voting (Alonso and Câmara (2016),Chan et al. (2019)), dynamic bank runs (Ely (2017)); stress testing (Inostroza and Pavan (2018)); auctions (Bergemann et al. (2017)); contests (Zhang and Zhou (2016)) etc.

[1]Specifically, this is what is known as *content-based* recommender systems.

1.3 Is more information always good? An Example

Let us consider an example where two objects, a and b, are to be allocated between two agents, 1 and 2, where each agent is supposed to be assigned exactly one object. Suppose the allocation mechanism used is Agent Proposing Deferred Acceptance (Gale and Shapley, 1962) where both the objects have higher priority for agent 1 than 2. Therefore this mechanism is equivalent to a Serial Dictatorship (SD) (Abdulkadiroğlu and Sönmez, 1998) where agent 1 is the dictator. Agents have i.i.d. preference distributions - the utility from object b is normalized to 0 and that from object a is distributed uniformly in $[-1, 1]$, so agents know ex-ante that they are equally likely to prefer a over b or b over a. Let u_{ij} denote agent i's utility from object j, $i \in \{1, 2\}, j \in \{a, b\}$.

Clearly, the expected utility from object a is equal to zero, which is the utility from object b. Hence, ex-ante both agents 1 and 2 are indifferent between the two objects. Under such an indifferent preference profile, depending on the tie-breaking rules of the mechanism, one of the two possible assignments is chosen, either $\{a \to 1, b \to 2\}$ or $\{b \to 1, a \to 2\}$, where $j \to i$ indicates object j is assigned to agent i, $i \in \{1, 2\}, j \in \{a, b\}$. No matter which one is chosen, the ex-ante aggregate welfare is $= \mathbb{E}(u_{1a} + u_{2b}) = \mathbb{E}(u_{2a} + u_{1b}) = 0$.

If agents are given full information about their own preferences, the four ordinal preference profiles and corresponding interim aggregate welfare outcomes which can arise under the SD mechanism are depicted in Figure 1. For example, in the bottom left square, where both agents have a negative utility from object a with probability one, and therefore prefer b over a, the allocation is $\{b \to 1, a \to 2\}$ under SD because 1 is the dictator. 1 gets a utility of zero from b, and 2 gets an average utility of -0.5 from a in this scenario. Therefore interim aggregate welfare for joint preferences lying in the bottom left square in the figure is -0.5. The corresponding quantities for the other three regions are calculated analogously. Due to the four regions being equally likely, the ex-ante aggregate welfare is $\frac{1}{4} \times (0.5 + 0.5 + 0.5 - 0.5) = \frac{1}{4}$.

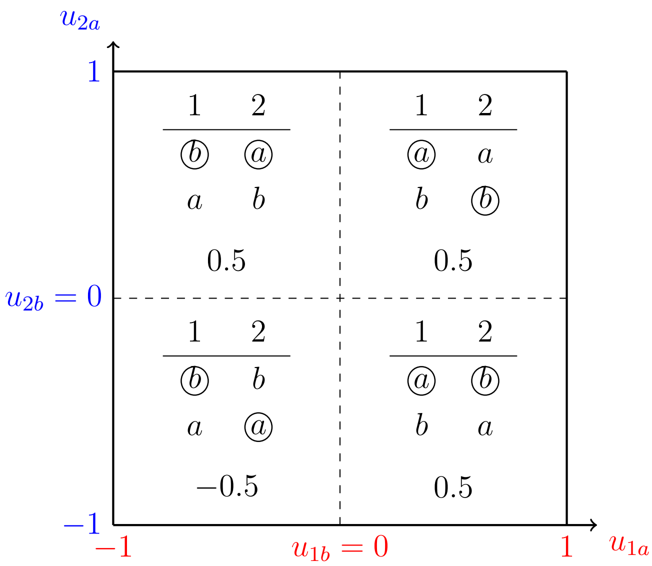

Figure 1: The X-axis (resp. Y-axis) plots agent 1's (resp. 2's) utility from object a, coded red (resp. blue). The four tables depict the four joint ordinal preference profiles which can arise, in their corresponding cardinal preference regions. The circled object names indicate the allocations under agent 1's dictatorship (SD) for each of these profiles. The numerical quantities at the bottom of each of the squares indicate the interim aggregate welfare conditional on the joint preferences lying in that region.

What does the planner ideally want? If possible, she wants to implement the point-wise welfare maximizing allocation at every realized preference profile. That is, if possible, she wants the allocation $\{a \to 1, b \to 2\}$ to be implemented whenever $u_{1a} + u_{2b} \geq u_{2a} + u_{1b}$ i.e. $u_{1a} \geq u_{2a}$ ($\because \ u_{1b} = u_{1b} = 0$), and the allocation $\{b \to 1, a \to 2\}$ to be implemented whenever $u_{1a} + u_{2b} \leq u_{2a} + u_{1b}$ i.e. $u_{1a} \leq u_{2a}$.

Now consider the "rank threshold" information policy indicated in Figure 2, which only reveals to the agents their *rank* in the empirical distribution of $\{u_{ia}\}_i$, i.e. only whether $u_{1a} \geq u_{2a}$ or $u_{2a} \geq u_{1a}$. When the dictator, agent 1, knows his rank is 1, his posterior expected value for a is $\mathbb{E}(u_{1a} | u_{1a} \geq u_{2a}; u_{1a}, u_{2a} \overset{i.i.d.}{\sim} U[-1, 1]) = 0.5 > 0$. Therefore *a posteriori* he prefers a over b, and hence picks a under SD, so the allocation is $\{a \to 1, b \to 2\}$ (The upper triangle in Figure 2). Analogously if agent 1's rank is 2 he prefers b in posterior, and picks b, so the allocation becomes $\{b \to 1, a \to 2\}$ (The lower triangle in Figure 2). Therefore this information policy achieves the maximum possible ex-ante aggregate welfare, which is strictly greater than that under full information: $\frac{1}{2} \times (0.5 + 0.5) = \frac{1}{2} > \frac{1}{4}$.

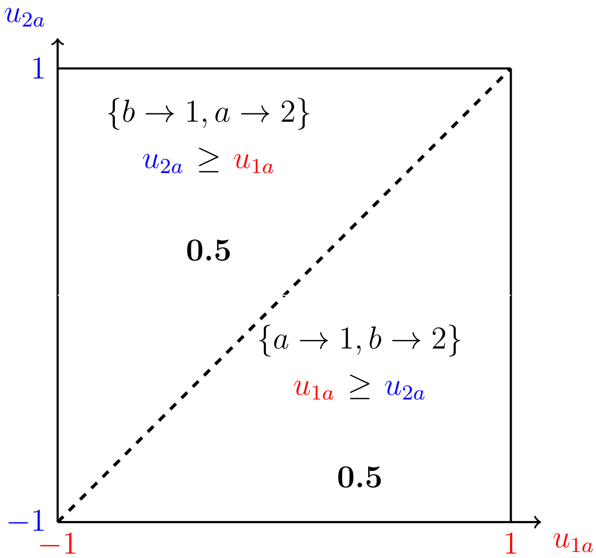

Figure 2: Instead of their respective ordinal preferences, agents are only informed if their joint preference lies in the upper (resp. lower) triangular region, where $u_{2a} \geq u_{1a}$ (resp. $u_{1a} \geq u_{2a}$). The resulting allocations under SD in each of the two regions are given in curly brackets. The numerical quantities in the bottom in each of the two triangular regions indicates interim aggregate welfare.

The above example might give the impression that while full information need not be optimal, we could expect it to at least improve welfare over the fully uninformative scenario. In section 3 we provide an example which shows this is not necessarily the case.

2 Model

2.1 Notation

All sets considered in this paper are subsets of topological spaces with their standard topologies, wherever applicable, and endowed with the Borel σ-algebra. For any such set X, we use $\mathcal{B}(X)$ to denote the Borel σ-algebra on X and ΔX to denote the set of all Borel probability measures over X, endowed with the weak-* topology. $\mathbb{1}_X(\cdot)$ denotes the indicator function indicating inclusion in the set X.

2.2 Setting

Ours is a standard model of object allocation without transfer, with cardinal preferences of agents (see Hylland and Zeckhauser (1979), more recently He et al. (2018) etc.), with the additional feature that agents only know the joint distribution of their cardinal preferences, rather than the preferences themselves. Formally, we consider an economy, $\Gamma = \{\mathcal{H}, I, Q, \mu\}$, where:

(i) $H = \{h\}_{h=1}^{m}$ is a set of objects.

(ii) $I = \{i\}_{i=1}^{n}$ is a set of agents (each referred to by *he*), each of whom is to be allocated exactly one object in total.

(iii) $Q = [q_h]_{h \in H}$ is a capacity vector, and $q_h \in \mathbb{N}$ is the number of copies of object $h, \forall h \in H$. We assume $\sum\limits_{h} q_h = n$, i.e. there is exactly one copy (of any object) for each agent. This is mainly for simplicity - almost all of our main results go through without this assumption, as discussed in Section 6.

(iv) Let $u_{i,h}$ denote agent i's utility from object h. $u_i \equiv$ agent i's utility vector. Each agent's i's utility vector lies in the compact set $\mathcal{U}_i \subset \mathbb{R}_+^m$. The joint utility space is $\mathcal{U} \equiv \times_i \mathcal{U}_i \subset \mathbb{R}_+^{m \times n}$. The typical element of $\mathcal{U}$ is denoted by u. u_is are distributed with joint prior distribution μ over $\mathcal{U}$, assumed atomless. μ is common knowledge.

The cardinal preference profile u - unknown to the agents - is sometimes referred to as the *state*.

In addition, objects may have a set of priorities (potentially weak) over the agents. We suppress this notation in the definition of the economy as this plays no role in our analysis.

Agents are expected utility maximizers. Let R_i represent agent i's ordinal ranking over objects, where P_i represents a strict ranking. Let $\mathcal{R}_i$ denote the set of all possible ordinal rankings over H. $R \equiv (R_i, R_{-i})$ denotes an ordinal preference profile of all agents. Let $\mathcal{R}$ denote the set of all possible ordinal preference profiles over H, with $\mathcal{P}$ denoting that of all possible *strict* ordinal preference profiles.

For all $u_i \in conv\ \mathcal{U}_i$ (resp. $u \in conv\ \mathcal{U}$) - where *conv* denotes the convex hull of any set - we use $R(u_i)$ (resp. $R(u)$) to denote the unique ordinal ranking (profile) consistent with the utility vector u_i (resp. profile u). Conversely, we use $\mathcal{U}_i(R_i)$ (resp. $\mathcal{U}(R)$) to denote the set of utility vectors (resp. profiles) consistent with the ordinal ranking (resp. profile) R_i (resp. R).

2.3 Mechanisms

Let $\mathcal{D}$ denote the set of all possible deterministic allocations of the copies of objects among the agents, subject to feasibility conditions. Each allocation $x \in \mathcal{D}$ (resp. $\Delta\mathcal{D}$) is represented by an $H \times I$ matrix with entries in $\{0, 1\}$ (resp. $[0, 1]$). The i-th column of x represents the allocation of agent i - i.e. x_{ij} gives agent i's probability of getting assigned object $j \in \{1, \cdots, m\}$ under allocation x. Hence feasibility conditions imply that for all $x \in \mathcal{D}$, $x_{ij} = \mathbb{1}_{\{i \text{ is allocated } j \text{ under allocation } x\}}$, and for all $x \in \Delta\mathcal{D}$, $\sum\limits_{j \in H} x_{ij} = 1\ \forall\ i$, $\sum\limits_{i \in I} x_{ih} = q_h\ \forall\ h$. For deterministic allocations $p \in \mathcal{D}$, we use $\widehat{p}_i \in H$ to denote the object received by agent i under the allocation p.

Let $\mathcal{M}_i$ be the space of allowable reports for agent i, which are used as inputs to the allocation mechanism. Let $\mathcal{M} = \times_i \mathcal{M}_i$ be the space of joint reports. A mechanism is a mapping $p : \mathcal{M} \to \Delta\mathcal{D}$, which maps joint reports of agents to distributions over outcomes. For example, for ordinal mechanisms[2] $\mathcal{M}_i = \mathcal{R}_i$, for cardinal ones $\mathcal{M}_i = \mathcal{U}_i$ etc.

We use $p_i(\mathcal{M}) \in \Delta H$ to denote the vector indicating the lottery over objects which agent i receives under mechanism p when the joint report is $m \in \mathcal{M}$. Analogously as above, if the allocation under mechanism p at joint report m is deterministic, i.e. $p(m) \in \mathcal{D}$, we use $\widehat{p}_i(m) \in H$ to denote the object received by agent i under the joint report m.

Wherever appropriate we use the notation $p(m) \cdot u$ (respectively $p_i(m) \cdot u_i$) to denote the aggregate (respectively agent i's) expected utility when the joint report is $m \in \mathcal{M}$ under allocation mechanism p, and the agents' utility profile is $u \in \mathcal{U}$. $\therefore\ p(m) \cdot u = \sum_{i \in I} p_i(m) \cdot u_i$.

In line with the convention in the literature, we call a mechanism *strategyproof* if it is dominant strategy incentive compatible.

We use (p, μ) to denote the *basic game* (Bergemann and Morris, 2016) - the set of possible reports $\mathcal{M}$, the common prior μ over $\mathcal{U}$ and the payoff functions - induced by the mechanism p under prior μ.

2.4 Principal's problem

The objective of the principal (she) is to maximize the society's aggregate welfare, as captured by the utilitarian welfare function - the sum of the expected utilities of the agents. This particular functional form is used only for ease of exposition. Most of our analysis and main results go through even when we replace the principal's objective function with any increasing function of the agents' interim expected utilities satisfying a mild condition, as discussed in Section 5.1.

We assume the mechanism $p \in \Delta\mathcal{D}^{\mathcal{M}}$ to be exogenously fixed. Therefore it is without loss of generality to restrict attention to *recommendation signals* (Bergemann and Morris, 2016; Taneva, 2019), i.e. signals which directly recommend agents to make specific reports to the mechanism. Therefore we assume, without loss of generality, that the signal space is $\mathcal{M}$.

A *signal* is a distribution over messages in $\mathcal{M}$ for each state u, chosen by the principal. Because this distribution depends on the state, it conveys information about the state, without necessarily fully revealing it. As an information designer, the goal of the principal is to use optimally chosen signals to steer agents towards joint actions she wants them to take.

In general signal realizations are assumed to be private - i.e. each agent observes

[2] Ordinal mechanisms are allocation mechanisms which use agents' reported ordinal rankings over objects as opposed to the cardinal strengths of their preferences over these objects.

only *his* recommendation and not those of others. We relax this assumption for particular cases in Section 4.

As mentioned before, we restrict attention to recommendation signals, so the incentive condition the principal's signal needs to satisfy is that - each agent must find it in his best interests to follow the principal's recommendations. This means, his expected utility from following his recommendation - conditional on the information his recommendation reveals to him *and* assuming each of the other agents follows his recommendation - must be weakly more than that from deviating. The constraints capturing this incentive notion are called *obedience constraints* (Bergemann and Morris, 2016).

For ease of exposition, we present the principal's problem for finite $\mathcal{M}$ here. The case of a general compact $\mathcal{M}$ is discussed in section 5.3.

The principal's problem is:

$$\max_{\substack{\nu(m|u):\\ \mathcal{U}\to\Delta\mathcal{M}}} \int_{u\in\mathcal{U}} \left(\sum_{m\in\mathcal{M}} (p(m)\cdot u)\,\nu(m|u)\right)\mu(du) \tag{P}$$

Subject to the obedience constraints:

For all $m_i, m_i' \in \mathcal{M}_i, i \in I$ such that $\nu(m_i) > 0$ and $p_i(m_i, m_{-i}) \neq p_i(m_i', m_{-i})$ for some m_{-i}:

$$\begin{aligned}
&\mathbb{E}_{(m,u)\sim\nu}\left[(p_i(m_i,m_{-i}) - p_i(m_i',m_{-i}))\cdot u_i|m_i\right] \geq 0,\\
&\text{i.e. } \mathbb{E}_{m_{-i}}\left[(p_i(m_i,m_{-i}) - p_i(m_i',m_{-i}))\cdot \mathbb{E}(u_i|(m_i,m_{-i}))|m_i\right] \geq 0,\\
&\text{i.e. } \sum_{m_{-i}}\left[(p_i(m_i,m_{-i}) - p_i(m_i',m_{-i}))\cdot\left(\int_u u_i\nu((m_i,m_{-i})|u)\mu(du)\right)\right] \geq 0.
\end{aligned}$$

where the joint distribution over $\mathcal{M}\times\mathcal{U}$ is given by $\nu(m,U) = \int_{u\in U} \nu(m|u)\mu(du)$ for all $m\in\mathcal{M}, U\in\mathcal{B}(\mathcal{U})$.

A straightforward application of the extreme value theorem shows that a solution to (P) exists.

Lemma 1. *A solution to the principal's problem* (P) *exists.*

The details of the proof are provided in Appendix A.1.

We assume principal-preferred equilibrium selection, i.e. we assume whenever an agent is indifferent between reporting a recommended message and a different message, he obeys by reporting the recommended message. Therefore if there exist i, m_i, m_i' such that $p_i(m_i, m_{-i}) = p_i(m_i', m_{-i})$ for all m_{-i}, agent i will never deviate via m_i' when recommended m_i. Therefore such pairs of messages do not impose any obedience constraint.

Going forward, "full information" refers to the informational setting where each agent has full information about his own preference, and no information about other agents' preferences.

Before solving the principal's problem we provide another example to illustrate the non-triviality of the information design problem in this context.

2.5 Is *some* information always better than no information? An Example

Fix a deterministic allocation $p \in \mathcal{D}$. Define $\mathcal{U}_p \equiv \{u \in \mathcal{U} : p \cdot u \geq p' \cdot u \; \forall \, p' \in \mathcal{D}\}$, the set of cardinal preference profiles for which p maximizes the utilitarian social welfare. Consider a prior μ such that *supp* $\mu = \mathcal{U}_p$. Let $\mathbb{E}_\mu(u) = u_0$, the agents' expected cardinal preference profile *a priori*. Let $R(u_0) = R_0$, the unique ordinal preference profile consistent with the cardinal preference profile u_0. The allocation p is social welfare maximizing for all $u \in \mathcal{U}_p$, therefore social welfare maximizing at u_0. Therefore p is Pareto efficient at R_0. Therefore there exists some strict ordering $\succ$ of the agents such that the corresponding Serial Dictatorship $SD_\succ$ satisfies $SD_\succ(R_0) = p$.[3] Define $\mathcal{R}_p \triangleq \{R \in \mathcal{R} : SD_\succ(R) = p\}$.

We can see that when the mechanism is $SD_\succ$ and the prior is μ, providing no information is optimal and produces First Best. Moreover, it produces strictly greater ex-ante utility than any equilibrium where any *relevant* information is revealed, i.e. any equilibrium where any joint report $R \in \mathcal{R} \setminus \mathcal{R}_p$ occurs with positive probability.[4]

It follows that there exists $\epsilon > 0$ such that for all priors μ such that $\mu(\mathcal{U}_p) > 1 - \epsilon$ (including full support priors), providing no information is strictly better than providing full information.

The above example, together with Example 1 from the introduction, suggest that information could either help or hurt, and even when it helps, more is not always better, thereby hinting at the non-triviality of the information design problem here.

3 When is full information optimal?

The most popularly used informational setting for the matching model is that of full information. It is also a sensible intuition that any loss of welfare arising out of the informational setting must be mainly due to agents not knowing their own preferences well enough. For example, Chen and He (2017) - who demonstrate that students overpay for information in school choice in a lab setting - echoes this intuition by suggesting that education authorities can improve welfare by "providing more information" to students, because this would save "socially wasteful costs of information acquisition". As demonstrated through the illustrative examples, this

[3] It is well-known that any Pareto efficient deterministic allocation can be implemented using some Serial Dictatorship. See, e.g. Abdulkadiroğlu and Sönmez (1998).

[4] The set of such equilibria is non-empty. The full information equilibrium is an example. Because *supp* $\mu = \mathcal{U}_p$, all ordinal preference profiles occur with positive probability.

intuition need not always be true. Our goal in this section is to determine how often this is the case, so that the information design problem is trivial.

When there is just one agent, and no preference divergence between the principal and the agent, it is intuitive to see that providing full information is optimal, because information cannot have any negative impact on anybody's payoffs (Kamenica and Gentzkow, 2011). Extending this intuition, if agents act in complete collusion - always maximizing the aggregate payoff to the group, *as if* they are just one agent - full information should still be optimal, for the same reason. Perhaps surprisingly however, it turns out that this is essentially the *only* situation in which full information is optimal. Whenever agents' preferences over outcomes diverge - i.e. outcomes that do not maximize the aggregate payoff occur with positive probability in equilibrium when agents are fully aware of their own preferences - revealing full information almost always hurts the overall welfare. We formalize this insight below, which leads to our first main result.

Because we focus on full information in this section, it is without loss to assume the message space to be $\mathcal{U}_i$ for each agent. Note that a necessary condition for full information to be optimal is that it is incentive feasible, i.e. it satisfies the obedience constraints of problem (P). It is easy to see that under full information, the obedience constraints boil down to standard Bayesian incentive compatibility constraints. Therefore we can focus on the truth-telling equilibrium under full information, and hence it is without loss to consider revelation mechanisms $p : \mathcal{U} \to \Delta\mathcal{D}$.

3.1 Finite utility space

First, we present our characterization for the case where the joint type space $\mathcal{U}$ is finite. This is sufficient to prove our main impossibility result, Theorem 1. We subsequently relax this assumption in subsection 5.3.

Below we define some properties useful for our characterization.

Definition 3.1 (Utilitarianism)**.** *A direct mechanism $p^* : \mathcal{U} \to \Delta\mathcal{D}$ is utilitarian if it satisfies $p^*(u) \cdot u \geq p \cdot u \; \forall \, p \in \Delta\mathcal{D}, \; \forall \, u \in \mathcal{U}$.*

As suggested by the name, a utilitarian mechanism maximizes the aggregate welfare in the economy at every realized cardinal preference profile. Below we weaken this property to what we call *weak utilitarianism*, which is crucial to our main characterization in this section.

Definition 3.2 (weak utilitarianism)**.** *We call a direct mechanism $p : \mathcal{U} \to \Delta\mathcal{D}$ weakly utilitarian if it satisfies $p(u) \cdot u \geq p(u') \cdot u \; \forall \, u' \in \mathcal{U}, \; \forall \, u \in \mathcal{U}$. We call it weakly utilitarian on the support of a prior μ if it satisfies $p(u) \cdot u \geq p(u') \cdot u \; \forall \, u' \in \mathcal{U}, \; \forall \, u \in supp \; \mu$.*

Like utilitarian mechanisms, weakly utilitarian mechanisms also maximize welfare at every utility profile, but only over their own range, i.e. over $p(\mathcal{U}) \subseteq \Delta\mathcal{D}$. Clearly, all utilitarian mechanisms are weakly utilitarian, but the converse is not true. For example, any constant mechanism - i.e. one of the form $p(u) = \overline{p} \; \forall \; u$ - is weakly utilitarian but not utilitarian.

As mentioned earlier, a necessary condition for full information to be optimal is that the mechanism is BIC w.r.t. the prior. Therefore for a given mechanism, the set of priors for which the full information structure can be optimal is a subset of the set of priors for which it is BIC.

Let $BIC^p \subseteq \Delta\mathcal{U}$ denote the set of priors with respect to which the mechanism p is BIC under full information[5], i.e.

$$\begin{aligned} BIC^p := \{\mu \in \Delta\mathcal{U} : &\sum_{u_{-i}} \left[(p_i(u_i, u_{-i}) - p_i(u_i', u_{-i})) \cdot u_i\right] \mu(u_{-i}|u_i) \geq 0 \\ &\forall \; u_i, u_i' \text{ s.t. } \mu(u_i), \mu(u_i') \neq 0.\} \end{aligned} \tag{1}$$

By finiteness of $\mathcal{U}$, BIC^p is a subset of a finite dimensional Euclidean space and therefore inherits its usual topology. So the interior of BIC^p (denoted $intBIC^p$) - the set of priors in BIC^p, each of which has some neighborhood in which every prior is in BIC^p - is well-defined.

The next proposition formalizes the idea introduced in the beginning of this section - for almost any prior with respect to which a mechanism is BIC, weak utilitarianism is both necessary and sufficient for full information to be utilitarian welfare maximizing.

Proposition 1. *Suppose $\mathcal{U}$ is finite. For any direct mechanism $p : \mathcal{U} \to \Delta\mathcal{D}$ such that $intBIC^p \neq \emptyset$, and any prior $\mu \in intBIC^p$, full information is utilitarian welfare-maximizing for the basic game (p, μ) if and only if p is weakly utilitarian on the support of μ.*

---

[5]Note that depending on p and $\mathcal{U}$, BIC^p may be empty. The following example illustrates this. Take any strategyproof mechanism $p^0 : \mathcal{U} \to \Delta\mathcal{D}$ under which there exists $i \in I, \overline{u}_i, \overline{u}_i' \in \mathcal{U}_i$ such that $p^0(\overline{u}_i, u_{-i}) \neq p^0(\overline{u}_i', u_{-i})$ for all u_{-i}. (For example, if p^0 is a serial dictatorship, agent i is the dictator, and each $u_i \in \mathcal{U}_i$ has a different, strictly most-preferred object.) Define the mechanism $p : \mathcal{U} \to \Delta\mathcal{D}$ by exchanging the allocations for the reports $(\overline{u}_i, u_{-i})$ and $(\overline{u}_i', u_{-i})$ for all u_{-i}. Formally:

$$p(u) = \begin{cases} p^0(u) & \text{if } u_i \notin \{\overline{u}_i, \overline{u}_i'\} \\ p^0(\overline{u}_i', u_{-i}) & \text{if } u_i = \overline{u}_i \\ p^0(\overline{u}_i, u_{-i}) & \text{if } u_i = \overline{u}_i' \end{cases}$$

p is clearly not BIC with respect to any prior $\mu \in \Delta\mathcal{U}$.

Moreover, even when BIC^p is non-empty, $intBIC^p$ may be empty. For example for any strategyproof p^0, and any $i \in I, \overline{u}_i, \overline{u}_i' \in \mathcal{U}_i$ such that $p^0(\overline{u}_i, u_{-i}) \neq p^0(\overline{u}_i', u_{-i})$ for some u_{-i}, p constructed from p^0 as above would have $intBIC^p = \emptyset$.

Because a weakly utilitarian p pointwise maximizes welfare at each $u \in supp\ \mu$ on $p(supp\ \mu)$, if full information is incentive feasible, i.e. $\mu \in BIC^p$, it is optimal. On the other hand, if $\mu \in intBIC^p$, none of the incentive constraints bind at the full information signal. So at every realized profile u, if with a small positive probability, a joint report in the support of $\underset{\widetilde{p} \in p(supp\ \mu)}{\arg\max}\ \widetilde{p} \cdot u$ is recommended - while u is revealed with the complementary (large) probability - it will be acceptable to a set of Bayes-rational agents. Combining these two insights, Proposition 1 follows. The details of the proof are provided in Appendix A.2.

Recall that if $intBIC^p \neq \emptyset$, $intBIC^p$ is a (Lebesgue) measure one subset of BIC^p. Therefore Proposition 1 tells us that, as long as full information is incentive-feasible, weak utilitarianism of the allocation mechanism on the support of the prior is not only sufficient, it is also almost always necessary, for full information to be welfare-maximizing.

Weak utilitarianism can also be thought of as a strong notion of incentive compatibility which resists collusion by the grand coalition of all agents, when they can exchange side payments.[6] Under the standard assumption of agents' utilities being quasilinear in money, if, for some joint preference profile, the coalition can jointly misreport in a way such that the sum of their expected utilities strictly increases, each agent can be made better off through such misreporting using side payments. Hence, under the aforementioned notion of incentive compatibility, the grand coalition always wants to maximize the sum of the expected utilities of all coalition members, i.e. all agents. Therefore this can be interpreted as the grand coalition acting *as if* they are one agent.

Clearly, for a benevolent principal designing information for one agent, revealing full information is optimal for any decision problem. The significance of Proposition 1 is that, it tells us that the converse also holds for almost all priors - full information is optimal only if the group of agents behave *as if* they are one agent, under the given allocation rule.

In the next section we focus on ordinal allocation mechanisms - highly popular in the literature and in practice - and apply the above characterization to answer the question - when is full information optimal for most standard ordinal object allocation mechanisms?

[6]This notion of incentive compatibility is related to, but weaker than, the incentive notion of *n-truthfulness*, which requires mechanisms to be resistant to collusion by any group of agents of any size who can exchange side payments. See, for example, Goldberg and Hartline (2005), Penna and Ventre (2008), Ji and Chen (2017) etc.

3.2 Ordinal mechanisms: An impossibility result

An immediate implication of Proposition 1 is the essential incompatibility of full information revelation with the utilitarian objective, in case of ordinal mechanisms. Clearly, if the prior has full support, any (non-constant) ordinal mechanism p is not weakly utilitarian. Therefore for the typical prior, full information is either infeasible as it violates incentive conditions, or strictly suboptimal. This is captured in Theorem 1 below. Recall that we endow the space of priors $\Delta\Theta$ with the weak-* topology.

Theorem 1. *There exists a generic set of priors $M \subset \Delta\Theta$ such that for any prior $\mu \in M$ and any non-constant ordinal mechanism $p : \mathcal{R} \to \Delta\mathcal{D}$, full information is not welfare-maximizing. Moreover, if p is strategyproof, then for any full support prior, full information is not welfare-maximizing.*

Note that for Theorem 1 we do not need to assume finiteness of $\mathcal{U}$.

The proof of Theorem 1 is provided in Appendix A.3.

Theorem 1 immediately leads to the following interesting corollary about some of the most popular ordinal mechanisms. Because of their strategyproofness, by Theorem 1, it is never socially optimal to reveal full ordinal information to agents, when any of these mechanisms is used.

Corollary 1. *When the allocation rule is fixed to be a serial dictatorship, deferred acceptance, random priority or top trading cycles, and the prior has full support, it is not welfare-maximizing to truthfully reveal each agent's ordinal preference ranking to him.*

4 When is first best achievable?

Clearly, the principal ideally wants to implement the pointwise maximum social welfare if possible. In this section we explore when and how that is possible. We define the first-best as the pointwise maximum aggregate social surplus that can be obtained.

Definition 4.1 (First Best)**.** *The First Best aggregate surplus is defined as:*

$$\int_u \max_{p \in \mathcal{D}} (p \cdot u)\mu(du)$$

4.1 The Object Recommendation signal

Next, we define a simple signal we call the *Object Recommendation signal (OR signal)*, $\nu_{OR} : \mathcal{U} \to \Delta\mathcal{D}$. As the name suggests, at each profile $u \in \mathcal{U}$, ν_{OR} *recommends* each agent to pick a particular object - i.e. rank it at the top of his preference ordering. The way it does this is by randomly choosing a utilitarian welfare-maximizing allocation, and then privately recommending each agent to pick the object he is assigned under this allocation. Formally,

$$\nu_{OR,i}(a|u) = \frac{|\{p \in \mathcal{D} : p \cdot u \geq p' \cdot u \; \forall \, p' \in \mathcal{D}, \widehat{p}_i(u) = a\}|}{|\{p \in \mathcal{D} : p \cdot u \geq p' \cdot u \; \forall \, p' \in \mathcal{D}\}|}, \; \forall \, i \in I, a \in H.$$

As the next characterization shows, the significance of the OR signal is that it characterizes the implementability of the first-best when the Deferred Acceptance algorithm is run without priorities, which is equivalent to Random Serial Dictatorship.

Theorem 2. *First Best is implementable by Deferred Acceptance without Priorities if and only if it is implementable by the OR signal.*

The details of the proof are provided in the appendix.

Armed with the above characterization, our goal is to identify the set of mechanisms and priors which allow us to implement the first-best. We begin by defining a class of mechanisms satisfying a mild efficiency criterion we call *weak efficiency*, which is simply the property that when there are no conflicts among the most preferred object demanded by each agent, everybody gets what he wants.

Definition 4.2 (Weak efficiency)**.** *A mechanism satisfies weak efficiency if under it, whenever the allocation assigning each agent (one of) his top reported object(s) is feasible, it is chosen.*

Obviously, efficient mechanisms such as Serial Dictatorship are weakly efficient, but the converse need not hold. For example, while Deferred Acceptance mechanisms are not efficient, they are *weakly efficient* in the above sense.

From the planner's point of view, the ideal situation is if at any true cardinal preference profile, a posterior ordinal preference profile can be realized (using some signal) where every agent wants the object he would be assigned under a utilitarian welfare maximizing allocation. Of course, in general, there need not exist *any* signal which makes this possible. For example, consider a setting where a priori, agents know their true ordinal preferences, under which everyone has the same most preferred object. In this case, obviously there is no signal which can give rise to the aforementioned class of posterior preference profiles. The class of priors for which it is always possible to create such a posterior preference profile is said to satisfy the *posterior top condition.*

Definition 4.3 (Posterior top condition)**.** *If the joint prior is such that, at every preference profile $u \in \mathcal{U}$, every agent's recommended object under the OR signal is (among) his posterior most preferred object(s), it is said to satisfy the posterior top condition.*

It is easy to see that if the posterior top condition is satisfied, First Best can always be implemented using any weakly efficient mechanism. This is formalized below.

Lemma 2 (Posterior top condition)**.** *If the posterior top condition is satisfied, then First Best is implementable by any weakly efficient mechanism.*

The proof is straightforward and is based on the observation that if the posterior top condition is satisfied, knowing that the mechanism is weakly efficient, agents know that if everyone obeys their recommendation, they would get the recommended object if they report it as their most-preferred. The details are provided in the Appendix.

An obvious but useful implication of Lemma 2 is that, when the Posterior top condition is satisfied, the Deferred Acceptance mechanism with *any* priority structure can implement First Best. This is significant because this mechanism is popular in several applied settings such as school choice, but is not efficient.

Corollary 2. *If the prior satisfies the posterior top condition, First Best is implementable by Deferred Acceptance with any priority structure.*

While it is obvious that the posterior top condition is sufficient for implementability of First Best by any weakly efficient mechanism, turns out, when the agents are symmetric, it is also necessary - in case of some popular weakly efficient mechanisms.

We say agents are *symmetric* if the distribution of their joint cardinal preference profile does not depend on the identities of the agents. Formally, if the prior μ is such that for any Borel subset $\widetilde{\mathcal{U}}$ of $\mathcal{U}$, and any permutation σ of $\{1, \cdots, n\}$, $\mu(u \in \widetilde{\mathcal{U}}) = \mu(u_\sigma \in \widetilde{\mathcal{U}})$, where u_σ is the vector obtained by permuting the components of u according to σ. Obviously, if the agents' preferences are i.i.d., they are symmetric. But they can be symmetric even when their preferences are not independent.

When the aforementioned symmetry property is satisfied, we obtain the following four-way equivalence.

Proposition 2. *Suppose agents are symmetric. Then the following are equivalent.*

1. *There exists an ordering of agents such that First Best is implementable by Serial Dictatorship.*

2. *First Best is implementable by Deferred Acceptance without priorities (i.e. Random Serial Dictatorship).*

3. *First Best is implementable by any weakly efficient mechanism.*

4. *The posterior top condition is satisfied.*

The proof is straightforward and is provided in the Appendix.

Next we provide sufficient conditions on the prior so that the posterior top condition is satisfied. To help drive intuition, we first note some interesting features of the OR signal when there are just two objects, and subsequently generalize them.

4.2 The case of two objects

Throughout this subsection we refer to the two objects as a and b.

Note that for two objects, agents' ordinal preferences are binary - object a (potentially weakly) preferred to object b or the reverse. Therefore, as long as we are using ordinal mechanisms, the only potential role any signal can play is to reverse an agent's relative preference over the two objects. We call a signal **informative on an agent's ordinal preferences**, if the posterior (weak) ordinal preference of the agent under the signal is the opposite of his prior preference.[7] As long as a signal can do that "whenever necessary", it can achieve the planner's most desired outcome - i.e. First Best. This insight is formalized in the proposition below.

Proposition 3. *Suppose $m = 2$ and the prior has full support. First best is achievable by any weakly efficient mechanism if and only if the OR signal is informative on each agent's ordinal preferences.*

Recall the example from Section 1.3 in the introduction, in which we introduced a *rankings* based information policy, one which conveyed to each agent only if his *relative preference* for object a over object b was higher or lower than that of the other agent - in other words, his *rank* in the empirical distribution of realized relative preferences. We saw in that example with two agents, that this signal achieved First Best. Turns out that in case of two objects, under certain conditions, this rank-based information policy is *equivalent to* the OR signal, and therefore its effectiveness generalizes to any number of agents as long as there are just two objects.

Without loss of generality, let us assume $q_a \geq q_b$. Let $\widehat{u}_{i,(a,b)} \equiv u_{ia} - u_{ib}$ denote agent i's relative preference for a over b. For a given utility profile $u \in \mathcal{U}$, let $rank(\widehat{u}_{i,(a,b)})$ denote the rank of agent i when all agents in I are sorted from highest to lowest according to $\widehat{u}_{i,(a,b)}$.

We formalize the rank threshold signal introduced in the example in Section 1.3 as follows. When there are just two objects, the rank threshold signal conveys to each agent whether his rank in the empirical distribution of each agent's relative preference for object a is greater than the number of copies of object a available. Formally:

[7]Therefore, by definition, if the agents are indifferent between the two objects a priori, as in Example 1, every signal is informative on each agent's ordinal preferences.

Definition 4.4 (Rank threshold signal)**.** *Suppose $m = 2$. At each realized cardinal preference profile $u \in \mathcal{U}$, the rank threshold signal informs each agent i whether $rank(\widehat{u}_{i,(a,b)}) > q_a$.*

As mentioned earlier, turns out that for i.i.d. preferences of agents and two objects, the rank threshold signal is equivalent to the OR signal.

Lemma 3. *Suppose $m = 2$ and agents' preferences are i.i.d.. Then the OR signal is equivalent to the rank threshold signal.*

The corollary below immediately follows from Lemma 3.

Corollary 3 (Rank threshold signal)**.** *Suppose $m = 2$ and agents' preferences are i.i.d.. First best is achievable by any weakly efficient mechanism if and only if the rank-threshold signal is informative on each agent's ordinal preferences.*

Corollary 3 tells us that as long as we are using reasonable mechanisms, the highest possible social welfare can be realized in a mechanism-agnostic way, if and only if informing agents of only their ranks in the distribution of relative preferences is decisive in *reversing* their ordinal preferences a posteriori.

However what makes the rank-based signal even more useful is that when agents are symmetric, regardless of whether First Best achieved, it *always* improves welfare over the case when agents are not provided any information at all. Hence if there are just two options, it is always useful to provide rank-related information to agents. This is formalized in the observation below.

Proposition 4. *Suppose $m = 2$ and agents are symmetric. Then the rank threshold signal weakly improves welfare over no information under any weakly efficient allocation mechanism.*

The significance of the equivalence of the OR and rank-threshold signal for the two-objects case lies in the fact that rank-based information policies arise naturally in many contexts. For example many standardized tests around the world provide only percentiles - equivalent of ranks - to students, and not scores. Proposition 4 says that if there are just two options - say, a specialized high school and an ordinary high school - and agents have identical preferences over these options (which is overwhelmingly true for the aforementioned example) this masking of the true score may actually be welfare improving (taking the test score as a proxy for the student's fit with the specialized school).

Before ending this section we introduce a notion of strong a priori preferences, which ties in some of the ideas introduced in this section and provides a segue to the a general version of the same notion which is key to the characterizations introduced in the next section for any finite number of objects.

As suggested by Proposition 3 and Corollary 3, the key to First Best being generally implementable is that the agents must be *sufficiently suggestible*. For example, if an agent knows a priori with probability one that a is better for him than b - even though he does not know for sure the cardinal strength of his preference - no signal can change his posterior preference to object b. This is an example of having *very strong* a priori preferences - a condition which renders the OR or rank threshold signal uninformative, and therefore First Best not implementable in a mechanism-agnostic way. This idea is formalized in the definition below.

Definition 4.5 (Strong a priori preference for $m = 2$). *Agent i is said to have a strong a priori preference for object a (b) if $\mathbb{E}(\widehat{u}_{i,(a,b)}|rank(\widehat{u}_{i,(a,b)}) \leq q_a) < 0$ ($\mathbb{E}(\widehat{u}_{i,(a,b)}|rank(\widehat{u}_{i,(a,b)}) > q_a) > 0$).*

This is clearly equivalent to the notion of informativeness of the rank threshold signal introduced earlier. Therefore, using Corollary 3 we can immediately conclude the following.

Corollary 4 (Rank threshold signal). *Suppose $m = 2$ and agents' preferences are i.i.d.. First best is achievable by any weakly efficient mechanism if and only if agents have no strong a priori preferences.*

The purpose of introducing Definition 4.5 and Corollary 4 is to build a bridge between the two-objects case and the general case. Turns out, the notion of *suggestibility* - as captured by having no strong a priori preferences - remains relevant to characterizing implementability of First Best even in the general case, as we shall shortly see in the next section.

The omitted proofs from this section are given in the appendix.

4.3 The general case

Before we generalize the concepts introduced above to the case of any finite number of objects, we need a mild regularity condition on the prior. The regularity condition captures the idea that any information about agent i's preferences over a subset of objects affects his valuation of those objects more than it affects his valuation of other objects.

Let $\mathcal{U}_{i,H'}$ denote the space of agent i's utility over objects in $H' \subset H$, i.e. $\mathcal{U}_{i,H'} \equiv [0,1]^{|H'|}$. Recall that $\mathcal{B}(\mathcal{U})$ and $\mathcal{B}(\mathcal{U}_{i,H'})$ denote the Borel σ-algebras over $\mathcal{U}$ and $\mathcal{U}_{i,H'}$ respectively. We define regularity as follows.

Definition 4.6 (Regularity). *We say the prior $\mu \in \Delta\mathcal{U}$ is regular if, for all $i, \mathcal{I} \in \mathcal{B}(\mathcal{U}), \mathcal{I}_{i,H'} \in \mathcal{B}(\mathcal{U}_{i,H'})$,*

$$|\mathbb{E}_\mu(u_{i,h'}|\mathcal{I}_{i,H'}, \mathcal{I}) - \mathbb{E}_\mu(u_{i,h'}|\mathcal{I})| \geq |\mathbb{E}_\mu(u_{i,h}|\mathcal{I}_{i,H'}, \mathcal{I}) - \mathbb{E}_\mu(u_{i,h}|\mathcal{I})| \qquad \text{(Reg)}$$

for all $h' \in H', h \in H \setminus H'$.

As we can see from the above definition, regularity is vacuously satisfied when there are just two objects.

$\widehat{u}_{i,(a,b)}$ and $rank(\widehat{u}_{i,(a,b)})$ are defined in the same way as in the previous subsection. Below we introduce a suitable generalization of *strong a priori preference* introduced in case of two objects.

Definition 4.7 (Strong a priori preference)**.** *Agent i is said to have a strong a priori preference for object $a \in H$, if there exists an object $b \in H$, such that under prior μ,*

$$\mathbb{E}_\mu(\widehat{u}_{i,(a,b)} | rank(\widehat{u}_{i,(a,b)}) > q_a) > 0$$

In other words, we say agent i has a strong a priori preference for object a relative to its supply if, even when he knows that his relative preference for object a over some other object b is not among the top q_a in the population, he still believes a is better for him than b.

The most extreme case of having no strong a priori preference for any object is, of course, ex-ante indifference.

Definition 4.8 (Ex-ante indifference)**.** *Agent i is said to be ex-ante indifferent if, for all $a, b \in H$,*

$$\mathbb{E}_\mu(\widehat{u}_{i,(a,b)}) = 0$$

We are now ready to formulate a partial characterization of the optimal signal for a class of priors which satisfy the condition of showing no strong a priori preference for any object.

Theorem 3. *Suppose the agents' preferences are independent, the prior is regular and no agent has a strong a priori preference for any object. Then, the OR signal, combined with any weakly efficient mechanism, produces the first-best allocation.*

Comparing with the two-objects case, we see that in the i.i.d. case, while the condition of no strong a priori preferences fully characterizes the implementability of First Best by any weakly efficient mechanism in case of two objects (Corollary 4), it still remains a sufficient condition when the number of objects is generalized.

The proof of Theorem 3 is based on the observation that under its conditions, the posterior top condition is satisfied, and therefore, by Lemma 2, the result follows. The details are provided in Appendix A.6.

5 Applications and Extensions

5.1 Other forms of planner's objective function

Suppose instead of maximizing the unweighted sum of agents' ex-ante expected utilities, the planner wants of maximize a weighted average of their utilities, where the weights are exogenously given. It is easy to see that all our results go through with suitable modifications, because all we have to do is redefine the state as the heterogenously scaled cardinal preference profile, and correspondingly redefine the domain of the allocation functions. This is briefly discussed below.

Suppose the agents' weights are $(w_1, \cdots, w_n)$ where $w_i \in \mathbb{R}_+ \ \forall \ i$. The only change that occurs in the planner's problem is that now her objective function becomes:

$$\max_{\substack{\nu(m|u): \\ \mathcal{U} \to \Delta\mathcal{M}}} \int\limits_{u \in \mathcal{U}} \left(\sum_{m \in \mathcal{M}} \left(\sum_i w_i p_i(m) \cdot u_i \right) \nu(m|u) \right) \mu(du), \tag{P'}$$

while the obedience constraints remain the same as in problem P.

In this context, we use the notion of *weighted* weak utilitarianism, which is an intuitive extension of the notion of weak utilitarianism introduced before.

Definition 5.1 (weighted weak utilitarianism)**.** *We call a direct mechanism $p : \mathcal{U} \to \Delta\mathcal{D}$ weighted weakly utilitarian with weights $(w_1, \cdots, w_n)$ if it satisfies $\sum\limits_i w_i p_i(u) \cdot u_i \geq \sum\limits_i w_i p_i(u') \cdot u \ \forall \ u' \in \mathcal{U}, \ \forall \ u \in \mathcal{U}$. We call it weighted weakly utilitarian with weights $(w_1, \cdots, w_n)$ on the support of a prior μ if it satisfies $\sum\limits_i w_i p_i(u) \cdot u_i \geq \sum\limits_i w_i p_i(u') \cdot u \ \forall \ u' \in \mathcal{U}, \ \forall \ u \in supp \ \mu$.*

Note that nothing changes on the incentive side, so the definition of BIC^p and $intBIC^p$ remain unchanged. Therefore we can modify Proposition 1 as follows, in this context.

Proposition 5. *Suppose $\mathcal{U}$ is finite. For any direct mechanism $p : \mathcal{U} \to \Delta\mathcal{D}$ such that $intBIC^p \neq \emptyset$, and any prior $\mu \in intBIC^p$, full information is a solution to (P') if and only if p is weighted weakly utilitarian with weights $(w_1, \cdots, w_n)$ on the support of μ.*

It is easy to see that nothing changes in the proof of Proposition 1 due to introducing weights, therefore Proposition 5 follows.

The definition of almost-sure weak utilitarianism introduced in section 5.3 and Proposition 6 can be modified analogously.

All of the positive results introduced in Section 4, such as Proposition 2, Theorem 2 and Theorem 3, also go through by - again - replacing each u_i by $w_i u_i$ in the characterization of the utilitarian rule, given in Claim A.6.1.

Extending this idea further, let us replace the planner's objective function with any increasing function of the agents' interim expected utilities which satisfies the condition that the allocation maximizing it is essentially unique. Formally, let the planner's objective function be $\psi : \mathbb{R}^n \to \mathbb{R}$, given by $\psi(p_1(m) \cdot u_1, \cdots, p_n(m) \cdot u_n)$, satisfying $\mu\left(\{u \in \mathcal{U} : |\arg\max_{p \in \Delta\mathcal{D}} \psi(p_1 \cdot u_1, \cdots, p_n \cdot u_n)| = 1\}\right) = 1$. It is easy to see that even in this case the characterization of the optimality of full information (modified analogously as shown above), and the impossibility results captured by Theorem 1 go through, as does the posterior top condition, Lemma 2. The only result which would not, in general, hold for general objective functions as described above, is Theorem 3. It is easy to see why - because the characterization of utilitarianism, as described by Claim A.6.1, would change for non-linear objective functions, which would lead to corresponding changes in the conditions on the prior which would give us implementability of First Best.

5.2 The case of independent priors

In many settings it is natural to assume that agents' preferences are independent. Note that the first part of Theorem 1 has no bite in this case, because the independent priors themselves form a null subset of the set of all priors. However, both Proposition 1 and Theorem 1 essentially go through even when we restrict our universe of priors to the set of independent priors, as we would show in this section.

Analogous to section 3.1, let us first consider the case where $\mathcal{U} = \times_i \mathcal{U}_i$ is finite. The definition of the set of independent priors w.r.t. which the mechanism p is BIC remains the same, except the universe of allowable priors is now restricted to the set of independent priors, $\Delta^{Ind}\mathcal{U} := \{\mu \in \Delta\mathcal{U} : \mu(u) = \prod_{i \in I} \mu(u_i)\}$ where $\mu(u_i) = \sum_{u_{-i} \in \mathcal{U}_{-i}} \mu(u_i, u_{-i})$.

$$BIC^p \equiv \{\mu \in \Delta^{Ind}\mathcal{U} : \sum_{u_{-i}} \left[(p_i(u_i, u_{-i}) - p_i(u'_i, u_{-i})) \cdot u_i\right] \mu(u_{-i}|u_i) \geq 0 \;\forall\, u'_i, \forall\, u_i \in supp\ \mu\} \tag{2}$$

as given in (1), except, due to independence, $\mu(u_{-i}|u_i) = \mu(u_{-i}) = \prod_{j \neq i} \mu(u_j)$, for all $i, u_i \in \mathcal{U}_i, u_{-i} \in \mathcal{U}_{-i}$. Hence the statement of Proposition 1 remains the same as well. We modify that of the first part of Theorem 1 as follows, in this environment.

Theorem 1´ . *Fix a non-constant ordinal mechanism $p : \mathcal{R} \to \Delta\mathcal{D}$. For almost all independent priors for which full information is incentive-feasible, it is not welfare-maximizing.*

The second part of Theorem 1 about strategyproof mechanisms, of course, remains the same.

The idea here is that, by the same reasoning as in the proof of Theorem 1, the set of independent priors for which full information may be optimal forms a subset of the set of all independent priors which is defined as the finite union of subsets defined by vanishing of polynomial functions of the prior (refer to equation (5) in the Appendix). Any set so defined forms a null subset of the space, the finite union of which, is another null subset.

5.3 Compact utility space

As shown by Theorem 1, the characterization of conditions for the optimality of full information provided in Proposition 1 is sufficient to draw useful insights for the case when the mechanism has a finite domain - such as ordinal mechanisms. However, it is less useful when the domain of the mechanism is uncountable. For example, cardinal mechanisms - direct revelation mechanisms which take agents' report of their utility vectors as inputs (Hylland and Zeckhauser, 1979; He et al., 2018) - constitute an important class of such mechanisms whenever the joint utility space $\mathcal{U}$ is compact. Hence we use this section to provide a characterization of the optimality of full information analogous to Proposition 1, for the case of a general compact utility space. Since we focus on direct mechanisms, as in Proposition 1, the message space coincides with the state space $\mathcal{U}$. Therefore as a first step we generalize the principal's problem (P) in the case of compact message spaces.

Let $\nu(\cdot|u)$ denote the conditional measure over $\mathcal{M}$ for each $u \in \mathcal{U}$ chosen by the principal, i.e. $\nu(M|u) = \int\limits_{m \in M} d\nu(m|u)$ for all $M \in \mathcal{B}(\mathcal{M})$. With a slight abuse of notation we also use $\nu(\cdot)$ to denote the unconditional measure induced over $\mathcal{M}$ by the prior μ and the aforementioned signal, i.e. $\nu(M) = \int\limits_{u \in \mathcal{U}} \nu(M|u) d\mu(u)$ for all $M \in \mathcal{B}(\mathcal{M})$.

Principal's problem for a compact $\mathcal{M}$:

$$\max_{\nu:\,\mathcal{U}\to\Delta\mathcal{M}} \int\limits_{u\in\mathcal{U}} \left(\int\limits_{m\in\mathcal{M}} (p(m)\cdot u)\, d\nu(m|u) \right) d\mu(u) \tag{P´´}$$

Subject to the constraint that players have an incentive to follow recommendations almost always (obedience):

$$\nu(m : \mathbb{E}\left[(p_i(m_i, m_{-i}) - p_i(m'_i, m_{-i})) \cdot u_i | m_i\right] \geq 0 \; \forall \; m'_i \in \mathcal{M}_i, \; \forall \; i \in I) = 1$$

where the expectation is taken over other agents' recommendations - m_{-i} - and agent i's own utility vector - u_i - conditional on the recommendation received by agent i - m_i.

In this new context, the appropriate way to define the set of priors with respect to which a given mechanism is BIC, is to use the slightly weaker notion of *almost sure BIC*. We call a direct mechanism $p : \mathcal{U} \to \Delta\mathcal{D}$ *almost surely BIC* if at

μ-almost all cardinal preference profiles u, no agent can be better off in expectation by misreporting. Let BIC^p denote the set of priors with respect to which the mechanism p is *almost surely* BIC. Formally, $BIC^p := \{\mu \in \Delta\mathcal{U} : \mu(u : \int_{u_{-i}} [(p_i(u_i, u_{-i}) - p_i(u_i', u_{-i})) \cdot u_i]\, \mu(du_{-i}|u_i) \geq 0\ \forall\ u_i' \in supp\ \mu,\ \forall\ i) = 1\}$.

Similarly, the appropriate analog of weak utilitarianism now becomes *almost sure* weak utilitarianism.

Definition 5.2 (Almost Sure Weak Utilitarianism)**.** *Fix a mechanism $p : \mathcal{U} \to \Delta\mathcal{D}$ and a prior μ. We call p μ-almost surely weakly utilitarian if $\mu\left(u : p(u) \cdot u \geq p(u') \cdot u\ \forall\ u' \in \mathcal{U}\right) = 1$.*

With this slight weakening of the key definitions from Section 3, we have the following characterization of the conditions for the optimality of full information, for compact $\mathcal{U}$. Note that BIC^p is a subset of $\Delta\mathcal{U}$, which is a compact set in the weak-* topology and inherits that topology. So, like in Section 3, the interior of BIC^p - the set of priors in BIC^p, each of which has some neighborhood in which every prior is in BIC^p - is well-defined.

Proposition 6. *For any direct mechanism $p : \mathcal{U} \to \Delta\mathcal{D}$ and any prior $\mu \in$ int BIC^p[8], full information is optimal for the basic game (p, μ) if and only if p is μ-almost surely weakly utilitarian.*

With these modified definitions, the proof of Proposition 6 is almost identical to that of Proposition 1. The intuition is as follows. Suppose at each joint utility profile the principal reveals full information to each agent with a high probability, but with the complimentary small probability she jointly recommends a report which would maximize social welfare at that profile. Because μ is interior to BIC^p, if the latter probability is small enough, each agent would find this recommendation acceptable always, even though he knows that the principal "lies" sometimes. As long as p is not μ-almost surely weakly utilitarian, there exists a positive μ-measure of joint utility profiles for which this would strictly increase aggregate ex-ante social welfare.

The details of the proof are provided in the Appendix.

5.4 Constrained utilitarian allocation

In many practical settings, the assumption of agents being ex-ante indifferent across all objects may not hold. So the first part of Theorem 3 would have no bite in such cases. However, in such settings, it is much more likely that such indifference would hold within a small *subset* of objects. For example, in the context of school

[8]All the caveats about *int* BIC^p mentioned in footnote 5 still apply, i.e. there exist mechanisms $p : \mathcal{U} \to \Delta\mathcal{D}$ for which *int* BIC^p is empty. However it is non-empty for most standard cardinal mechanisms described in the literature. See, e.g. Hylland and Zeckhauser (1979), He et al. (2018) etc.

choice, where all the schools in a school district are in the choice set of students, it may not be realistic to assume that a student is indifferent between two schools which are academically comparable but one of them is significantly further away from their home compared to the other. However, it is much more likely that the same student would be ex-ante indifferent among schools which are within a few blocks from their home.

Such limited ex-ante indifference might still be useful for practical purposes when considered together with the fact that in many settings there are restrictions on the set of objects each agent is eligible for. For example, in the school choice context, often the eligibility of each student is restricted to schools nearby him (see, e.g. Shi (2016)). The relevant welfarist question here is, what is the maximum possible aggregate social welfare that can be achieved *with these restrictions*? We call this the *constrained first-best*. As is intuitive, each agent being ex-ante indifferent across objects he is eligible for is sufficient for achieving it in such settings, when other conditions from Theorem 3 are met, as formalized below.

Corollary 5. *Suppose agent preferences are independent, the prior is regular, and each agent is ex-ante indifferent within within the set of objects he is eligible for. Then, the OR signal combined with any weakly efficient mechanism produces a constrained first-best allocation.*

5.5 Controlling both the mechanism and the information structure

A natural question that arises in the settings of object allocation we have studied is: What if the designer can design both the mechanism and the information structure?

In such a setting the principal's problem becomes:

$$\max_{M=\times_i M_i}\left[\max_{p:M\to\Delta\mathcal{D}}\left[\max_{\substack{\nu:\\ \mathcal{U}\to\Delta M}}\int_{u\in\mathcal{U}}\left(\int_{m\in\mathcal{M}}(p(m)\cdot u)\,d\nu(m|u)\right)\mu(du)\right]\right]$$

Subject to obedience constraints:

$$\mathbb{E}\left[(p_i(m_i,m_{-i})-p_i(m_i',m_{-i}))\cdot u_i|m_i\right]\geq 0\;\forall\; m_i,m_i'\in M_i,\;\forall\; i\in I$$

Note that without additional constraints, this is a trivial problem. For example, consider the following mechanism: Each agent can ask for exactly one object. If the allocation when each agent is given the object he asked for is feasible, it is implemented. If it is not feasible, no object is allocated to any agent. It is easy to see that this mechanism combined with the (public or private) OR signal always achieves first best, even when $\sum_h q_h \geq n$.[9] However, it may be interesting to explore its implications if there are additional constraints (e.g. agents' signals must be independent).

6 Discussion

6.1 Bypassing the impossibility

The full support assumption in the second part of Theorem 1 is important. Without it, there are several important exceptions for which full information is indeed optimal, as we illustrate below.

First, consider the obvious case of n agents and n objects, $h_1, \cdots, h_n$ such that $\mu(a_i P_i a_j \ \forall\ i \neq j) = 1$, i.e. every agent has a commonly known, different top object. In this case full information is obviously optimal. In fact it achieves first-best.

A simple but practically relevant domain of preferences in the context of the assignment problem is the **dichotomous preference domain** - a domain where each agent views each object as either acceptable or unacceptable (Bogomolnaia and Moulin, 2004; Bogomolnaia et al., 2005). So the utility each agent gets from each object can be represented as either 0 or 1. Turns out, in such a domain, when there are just two objects (with potentially multiple copies) full information is optimal and, in fact, achieves first-best for a popular class of mechanisms.

Proposition 7 (Dichotomous domain)**.** *In any dichotomous preference domain where there are only two objects, full information together with any efficient and strategyproof mechanism is optimal and achieves first best.*

Strategyproofness in the statement of Proposition 7 can be weakened to BIC without altering the conclusions. Proposition 7 still uses strategyproofness to make it clear that for some of the popular mechanisms in the literature such as serial dictatorship and probabilistic serial which are efficient and strategyproof on the dichotomous domain[10], full information is welfare-maximizing whenever there are just two objects.

Another simple and intuitive domain of preferences is an **ordinal rank-value preference domain**. We call a domain of utilities $\mathcal{U}_i$ an *ordinal domain* if it contains at most one utility vector compatible with every possible ordinal ranking.

[9] In case of $\sum_h q_h \geq n$, when an agent receives a recommendation for an object, he knows that conditional on other agents accepting their recommendations - which is social welfare maximizing - if there existed an object other than the one recommended to him under the utilitarian allocation which would give him higher utility, he would have been allocated that object under the utilitarian rule, and would therefore have been recommended that object.

[10] The probabilistic serial is not strategyproof in general, but is so on the dichotomous domain. See Liu (2017).

An *ordinal rank-value domain* (Featherstone, 2011, 2014) is an ordinal domain with strict rankings where each rank is associated with a utility value, regardless of the object. For example, suppose $H = \{a, b, c\}$. The set of rank value domains is given by the set of $3! = 6$ permutations of any three distinct non-negative numbers.

As is apparent, rank value domains are characterized by their symmetry. This symmetry property is further enhanced in settings where the possible utility vectors are not only permutations of each other, but also equally likely. Turns out the symmetry induced by these two factors taken together is sufficient for any utilitarian mechanism to be BIC, which means full information is welfare-maximizing for such mechanisms in this setting. This is formalized below.

Proposition 8. *When the utility space $\mathcal{U}$ satisfies the ordinal rank-value property, under an independent uniform prior, for any utilitarian mechanism, full information is optimal and achieves First Best.*

The proofs are provided in the appendix.

6.2 Regularity

The regularity assumption in Section 4.3 is stronger than what is needed for Theorem 3 to hold. We defined regularity as we did in that section because it is more intuitive and we believe, more generally relatable. However, all we need for Theorem 3 is that "good news" about an object increases its posterior expected value to any agent more than it increases - if at all - the posterior expected value of other objects to that agent. This is sufficient because when the OR signal recommends an object to an agent, it conveys positive news about that object to the agent, in a sense we will make precise below.

We use notation from Section 4. Fix an object $a \in H$. Let $H' \subseteq H \setminus a$ and $v_{ah} \in [-1, 1]$ for all $h \in H'$. We call information sets of the form $\bigcap\limits_{h \in H'} \{\hat{u}_{i,ah} \geq v_{ah}, v_{ah}\}$ *positive news* about object a to agent i, as they convey that agent i's payoff from object a is greater than that from each of the objects in H' by some additive factor (positive or negative). This allows us to weaken regularity to the notion of *regularity** defined below.

Definition 6.1 (Regularity*). *We say the prior $\mu \in \Delta\mathcal{U}$ is regular* if, for all $i \in I, a, b \in H$ and all positive news about a to i of the form $\bigcap\limits_{c \in H \setminus \{a,b\}} \{\hat{u}_{i,ac} \geq v_{ac}, v_{ac}\}$, we have:*

$$\begin{aligned}
&\mathbb{E}_\mu(u_{i,a} | \hat{u}_{i,ab} \geq v_{ab}, \hat{u}_{i,ac} \geq v_{ac} \;\forall\, c \in H \setminus \{a, b\}) - \mathbb{E}_\mu(u_{i,a} | \hat{u}_{i,ab} \geq v_{ab}) \\
&\qquad \geq \mathbb{E}_\mu(u_{i,b} | \hat{u}_{i,ab} \geq v_{ab}, \hat{u}_{i,ac} \geq v_{ac} \;\forall\, c \in H \setminus \{a, b\}) - \mathbb{E}_\mu(u_{i,b} | \hat{u}_{i,ab} \geq v_{ab}).
\end{aligned}$$

Reularity* requires that when i already has positive news about object a relative to object b (i.e. i's information set is $\{\hat{u}_{i,ab} \geq v_{ab}, v_{ab}\}$ for some $v_{ab} \in [-1, 1]$), *additional* positive news about a relative to all other objects increases i's posterior expectation of his value for a more than it increases his posterior expectation of his value for b.

It is clear from the proof of Theorem 3 that we can replace Regularity by Regularity* in Theorem 3 and it will still hold.

6.3 Number of copies of objects

It is clear from the proof of our main result, Theorem 1 and that of the first part (ex-ante indifference) of Theorem 3, that the assumption $\sum_h q_h = n$ can be relaxed to its most general alternative - $\sum_h q_h \geq n$. The only main result for which the assumption $\sum_h q_h = n$ is required is the second part of Theorem 3 - for the OR signal to deliver first-best when agent preferences are i.i.d. without any strong a priori preferences. In addition, relaxing the assumption $\sum_h q_h = n$ would change some of the results in Section 6.1.

7 Conclusion

In the context of the object allocation problem, this paper studies a scenario where agents are unsure about their preferences over objects and a central planner wants to design an *experiment* - such as an academic or professional or eligibility test - to strategically reveal information to agents about their own preferences, with the objective of maximizing ex-ante social welfare. I first attempt to make a case for information design in this setting, by showing that when the mechanism is ordinal, informing each agent perfectly about his own preferences is almost always strictly suboptimal. As a partial answer to the obvious next question, I then characterize the optimal information policy for a class of priors which satisfy the condition of not having too strong *a priori* preferences over objects. I show that for this class of priors, the optimal information policy achieves first-best - i.e. it is able to implement an allocation which maximizes social welfare at each possible realized preference profile. This optimal information policy turns out to be to simply recommend each agent to pick the object he would have received under the pointwise social welfare maximizing allocation.

To the best of my knowledge, this is the first attempt at analyzing the object allocation problem through the lens of information design. Therefore several areas of future research suggest themselves. First, in this model we assume agents have no private information to start with. This could to be relaxed to study a more general setting where agents have some private information about their own preferences. Secondly, we have allowed the signals of agents to be correlated - i.e. for the principal to provide *some* information about the preferences of one agent to another agent.

What if that is not allowed, for legal or practical reasons? Third, while without additional constraints the problem of jointly designing information and mechanism becomes trivial - as discussed in section 5.5 - this need not be the case with additional constraints. For example, the aforementioned independence requirement for agent signals, or other constraints on the mechanism (e.g. stability, fairness etc.) may restore non-triviality to the joint information and mechanism design problem. It would be interesting to study these important questions in future research.

Appendix

A.1 Proof of Lemma 1

The principal's problem (P) can be recast as that of maximizing a linear objective function over $\Delta(\mathcal{M}\times\mathcal{U})$ subject to linear equality (Bayes-plausibility) and inequality (obedience) constraints, as follows.

Define $V : \mathcal{M}\times\mathcal{U} \to \mathbb{R}$ as $V(m,u) = p(m)\cdot u$ for all m, u. Then (P) can be written as:

$$\max_{\nu\in\Delta(\mathcal{M}\times\mathcal{U})} \int_{(m,u)\in\mathcal{M}\times\mathcal{U}} V(m,u)\nu(m,du)$$

subject to:

A. Bayes-plausibility: $\int_U \sum_{m\in\mathcal{M}} \nu(m,du) = \mu(U)\ \forall\ U\in\mathcal{B}(\mathcal{U})$, and (A)

B. Obedience: $\nu(m : \mathbb{E}\left[(p_i(m_i,m_{-i}) - p_i(m_i',m_{-i}))\cdot u_i|m_i\right] \geq 0\ \forall\ m_i'\in\mathcal{M}_i,\ \forall\ i\in I) = 1.$ (B)

Therefore the feasible set is $\mathcal{F}(p,\mu) = \{\nu\in\Delta(\mathcal{M}\times\mathcal{U}) : \nu \text{ satisfies (A) and (B)}\}$. Below we show that $\mathcal{F}$ is compact.

It is a standard result that the subset of $\Delta(\mathcal{M}\times\mathcal{U})$ satisfying (A) - call this $\Delta^\mu(\mathcal{M}\times\mathcal{U})$ - is a compact set in the weak-* topology on $\Delta(\mathcal{M}\times\mathcal{U})$ (Khan et al. (2013), Section 2). Next we show that the subset of $\Delta^\mu(\mathcal{M}\times\mathcal{U})$ satisfying (B) is closed. Take a sequence of measures $\{\nu_n\}\in\Delta^\mu(\mathcal{M}\times\mathcal{U})$ such that $\nu_n \implies \nu$. $\therefore\ \nu\in\Delta^\mu(\mathcal{M}\times\mathcal{U})$. Fix i, m_i, m_i'. For any $\hat{\nu}\in\Delta(\mathcal{M}\times\mathcal{U})$ define the finite measure $\widehat{\nu}^{m_i}$ over the measurable space $(\mathcal{M}_{-i}\times\mathcal{U}, \mathcal{B}(\mathcal{M}_{-i}\times\mathcal{U}))$ as $\widehat{\nu}^{m_i}(X) = \hat{\nu}(\{m_i\}\times X)\ \forall\ X\in\mathcal{B}(\mathcal{M}_{-i}\times\mathcal{U})$. Therefore the sequence of measures $\{\nu_n^{m_i}\}$ is bounded above by 1 and $\therefore\ \nu_n^{m_i} \implies \nu^{m_i}$. $(p_i(m_i,m_{-i}) - p_i(m_i',m_{-i})\cdot u_i$ is a continuous function of (m_{-i},u), bounded by $\overline{u}$. Therefore by the Portmanteau theorem (Klenke, 2013),

$$\int\limits_{m_{-i},u} [(p_i(m_i, m_{-i}) - p_i(m_i', m_{-i})) \cdot u_i]\nu_n^{m_i}(m_{-i}, du)$$
$$\to \int\limits_{m_{-i},u} [(p_i(m_i, m_{-i}) - p_i(m_i', m_{-i})) \cdot u_i]\nu^{m_i}(m_{-i}, du)$$

$\{\int\limits_{m_{-i},u} [(p_i(m_i, m_{-i}) - p_i(m_i', m_{-i})) \cdot u_i]\nu_n^{m_i}(m_{-i}, du)\}_n$ is a sequence of real numbers. Therefore, $\int\limits_{m_{-i},u} [(p_i(m_i, m_{-i}) - p_i(m_i', m_{-i})) \cdot u_i]\nu_n^{m_i}(m_{-i}, du) \geq 0 \; \forall \; n \implies \int\limits_{m_{-i},u} [(p_i(m_i, m_{-i}) - p_i(m_i', m_{-i})) \cdot u_i]\nu^{m_i}(m_{-i}, du) \geq 0$. Therefore $\mathcal{F}$ is a closed subset of a compact set, and is therefore compact.

The mechanism p - and therefore the function V - is fixed. So the function $\mathbb{V} : \Delta(\mathcal{M} \times \mathcal{U}) \to [n\underline{u}, n\overline{u}]$ defined as the maximand, $\mathbb{V}(\nu) = \int\limits_{(m,u) \in \mathcal{M} \times \mathcal{U}} V(m, u)\nu(m, du)$, is bounded and linear - and therefore continuous - in ν, which is to be maximized over the compact set $\mathcal{F}$. Hence existence follows by the extreme value theorem.

A.2 Proof of Proposition 1

For finite $\mathcal{U}$, the principal's problem (P) becomes:

$$\max_{\substack{\nu(\widehat{u}|u): \mathcal{U} \times \mathcal{U} \to [0,1], \\ \sum\limits_{\widehat{u} \in \mathcal{U}} \nu(\widehat{u}|u) = 1 \; \forall \; u \in \mathcal{U}}} \sum_{\widehat{u}, u} (p(\widehat{u}) \cdot u) \, \nu(\widehat{u}|u)\mu(u) \tag{OptBCE}$$

Subject to the obedience constraints:

$$\sum_{\widehat{u}_{-i}, u} \mu(u)\nu\left((\widehat{u}_i, \widehat{u}_{-i}) \,|u\right) p_i(\widehat{u}_i, \widehat{u}_{-i}) \cdot u_i \geq \sum_{\widehat{u}_{-i}, u} \mu(u)\nu\left((\widehat{u}_i, \widehat{u}_{-i}) \,|u\right) p_i(\widehat{u}_i', \widehat{u}_{-i}) \cdot u_i, \tag{OBED}$$

for all $\widehat{u}_i, \widehat{u}_i' \in \mathcal{U}_i, i \in I$.

The full information signal is:

$$\nu^f(\hat{u}|u) = \begin{cases} 1, \text{for } \hat{u} = u \\ 0, \text{ otherwise} \end{cases}$$

First, note the following fact from linear programming which we also prove for the sake of completeness.

Fact 1. *If a solution to a linear program becomes infeasible when a finite set of additional linear inequality constraints are added, at least one of the new constraints must bind at the new solution, if it exists and is not a solution to the original problem.*[11]

[11] I thank Misha Lavrov for suggesting this line of argument.

Proof. Let $A^i \in \mathbb{R}^{l_i \times k}, b^i \in \mathbb{R}^{l_i}, i \in \{1,2\}, c \in \mathbb{R}^k, k, l_1, l_2 \in \mathbb{N}$. Consider the following two linear programs:

$$\max_{x \in \mathbb{R}^k} c^T x, \qquad \text{(LP1)} \qquad\qquad \max_{x \in \mathbb{R}^k} c^T x, \qquad \text{(LP2)}$$
$$\text{s.t. } A^1 x \leq b^1 \qquad\qquad \text{s.t. } \left[A^{1^T}\ A^{2^T}\right]^T x \leq \left[b^{1^T}\ b^{2^T}\right]^T$$

Suppose x^1 is a solution to (LP1), but ceases to be feasible for (LP2). Suppose x^2 is a solution to (LP2) but not to (LP1), and $A^2 x^2 < b^2$. Therefore for small enough $\epsilon \in (0,1)$, $\left[A^{1^T}\ A^{2^T}\right]^T (\epsilon x^1 + (1-\epsilon)x^2) \leq \left[b^{1^T}\ b^{2^T}\right]^T$, i.e. $(\epsilon x^1 + (1-\epsilon)x^2)$ is feasible for (LP2). x^2 is not a solution to (LP1), hence $c^T x^1 > c^T x^2$. $\therefore$ $c^T(\epsilon x^1 + (1-\epsilon)x^2) > c^T x^2$ and $(\epsilon x^1 + (1-\epsilon)x^2)$ is feasible for (LP2). Therefore x^2 cannot be a solution to (LP2), which is a contradiction. □

Now we are ready to prove the Proposition.

Proof. ($\Rightarrow$) If the mechanism is weakly utilitarian, the full information signal is a solution to the unconstrained problem (OptBCE). If it is BIC with respect to full information, that means the full information signal is in the feasible set even with the obedience constraints, and is therefore the constrained optimum.

($\Leftarrow$) Any optimal solution to the unconstrained problem (OptBCE) picks a distribution over the set of joint reports, each of which maximizes the total ex-ante utility at *each* joint utility profile. Suppose the mechanism is not weakly utilitarian. This means the full information signal is *not* an optimal solution to the unconstrained problem (OptBCE). But it is an extreme point of the unconstrained feasible set. Under these conditions, by Fact 1, the full information signal can become a constrained optimal BCE only if the prior distribution is such that at least one of the obedience constraints (OBED) binds at the full information signal, which is equivalent to the latter lying on the boundary of BIC^p defined by:

$$\begin{aligned} bd\ BIC^p &= \bigcup_{u_i, u_i', i} \{\mu \in \Delta\mathcal{U} : \mathbb{E}\left(\left(\left(p_i(u_i, u_{-i}) - p_i(u_i', u_{-i})\right) \cdot u_i\right) | u_i\right) = 0\} \\ &= \bigcup_{u_i, u_i', i} \{\mu \in \Delta\mathcal{U} : \sum_{u_{-i} \in \mathcal{U}_{-i}} \left(\left(p_i(u_i, u_{-i}) - p_i(u_i', u_{-i})\right) \cdot u_i\right) \mu(u_i, u_{-i}) = 0\} \end{aligned} \tag{3}$$

Therefore if $\mu \in int\ BIC^p$ and p is not weakly utilitarian on the support of μ, full information cannot be optimal. □

A.3 Proofs for Section 3.2

Proof of Theorem 1. Fix an ordinal mechanism $p : \mathcal{R} \to \Delta\mathcal{D}$.

Let π_{ord} be the signal corresponding to the (essentially unique) ordinal partition of $\mathcal{U}$, $\{\mathcal{U}(R)\}_{R \in \mathcal{R}}$. Formally, $\pi_{ord} : \mathcal{U} \to \Delta\mathcal{R}$, $\pi_{ord}(R|u) = \mathbb{1}_{\{u \in \mathcal{U}(R)\}}\ \forall\, u \in \mathcal{U}, R \in \mathcal{R}$.

We use π^+_{ord} to generically denote the signal corresponding to any finite refinement of the ordinal partition $\{\mathcal{U}(R)\}_{R\in\mathcal{R}}$. Formally, π^+_{ord} is a signal $\pi^+_{ord} : \mathcal{U} \to \Delta\mathcal{T}$, $\pi^+_{ord}(T|u) = \mathbb{1}_{\{u\in\mathcal{U}_T\}}\ \forall\ u \in \mathcal{U}, T \in \mathcal{T}$ where $\mathcal{T}$ is some indexing set such that $\{\mathcal{U}_T\}_{T\in\mathcal{T}}$ is a partition of $\mathcal{U}$, $|\mathcal{R}| \leq |\mathcal{T}| < \infty$ and $\mathcal{U}_T \subseteq \mathcal{U}(R)$ for some R for all $T \in \mathcal{T}$. Going forward, abusing notation, we use π_{ord} and π^+_{ord} to denote both the aforementioned signals and the partitions of $\{\mathcal{U}(R)\}_{R\in\mathcal{R}}$ they correspond to. Let Π^+_{ord} denote the set of all signals corresponding to finite refinements of the partition $\{\mathcal{U}(R)\}_{R\in\mathcal{R}}$, including π_{ord}. Let $\mu_{\pi^+_{ord}}$ denote the posterior belief over $\mathcal{U}$ induced by $\pi^+_{ord} \in \Pi^+_{ord}$. That is, for $\pi^+_{ord} \in \Pi^+_{ord}$ defined by $\pi^+_{ord}(T|u) = \mathbb{1}_{\{u\in\mathcal{U}_T\}}\ \forall\ u \in \mathcal{U}, T \in \mathcal{T}$, $supp\ \mu_{\pi^+_{ord}} = \{\mathbb{E}(u|u \in \mathcal{U}_T)\}_{T\in\mathcal{T}}$, with $\mu_{\pi^+_{ord}}(\mathbb{E}(u|u \in \mathcal{U}_T)) = \mu(\mathcal{U}_T)$, $\forall\ T \in \mathcal{T}$.

We are now ready to prove the theorem, which we do in the following steps.

Lemma A.3.1. *Fix any $\pi^+_{ord} \in \Pi^+_{ord}$. The following are equivalent:*

1. *The allocation rule p is BIC with respect to $\mu_{\pi^+_{ord}}$.*
2. *The full information signal is a BCE of the basic game $(p, \mu_{\pi^+_{ord}})$.*

Proof. Straightforward algebra show that, for any $\pi^+_{ord} \in \Pi^+_{ord}$, both BIC constraints and obedience constraints defined by equation (OBED) with respect to the full information signal are given by:

$\mathbb{E}\left(\left((p_i(R_i, R_{-i}) - p_i(R'_i, R_{-i})\right) \cdot \mathbb{E}(u_i|(R_i, R_{-i}))\right)|R_i\right) \geq 0,\ \forall\ i, R_i, R'_i.$

The claim follows. □

Lemma A.3.2. *π_{ord} is the optimal signal for p only if the full information signal induces an optimal BCE in the basic game $(p, \mu_{\pi^+_{ord}})$ for any $\pi^+_{ord} \in \Pi^+_{ord}$.*

Proof. p is BIC with respect to the prior μ, therefore BIC with respect to $\mu_{\pi^+_{ord}}$ for any $\pi^+_{ord} \in \Pi^+_{ord}$. Therefore by Lemma A.3.1, the full information signal induces a BCE of the basic game $(p, \mu_{\pi^+_{ord}})$ for all $\pi^+_{ord} \in \Pi^+_{ord}$.

Using $\nu^f : \mathcal{U} \to \Delta\mathcal{R}$ to denote the full information signal

Clearly, from the objective in (OptBCE), for the basic game $(p, \mu_{\pi^+_{ord}})$, $\mathbb{E}_{\nu^f}[V] = \sum_R \left(\sum_i p_i(R) \cdot \mathbb{E}_\mu(u_i|R)\right) \mu\left(\mathcal{U}(R)\right)$, which is not a function of π^+_{ord}, as long as $\pi^+_{ord} \in \Pi^+_{ord}$.

Suppose there exists a $\pi^+_{ord} \in \Pi^+_{ord}$ such that ν^f is not an optimal BCE of the basic game $(p, \mu_{\pi^+_{ord}})$. Then the optimal BCE of $(p, \mu_{\pi^+_{ord}})$, say ν, produces strictly higher ex-ante value than ν^f, i.e. $\mathbb{E}_\nu[V] > \mathbb{E}_{\nu^f}[V]$. But $\mathbb{E}_{\nu^f}[V]$ is the value for p when the information structure is π_{ord}. Therefore if such a $\pi^+_{ord} \in \Pi^+_{ord}$ exists then π_{ord} cannot be the optimal information structure for p. □

Define $\mathcal{U}^*(p) = \bigcup_{R\in\mathcal{R}} \{u \in \mathcal{U} : u \in \mathcal{U}(R), \sum_i p_i(R) \cdot u_i \geq \sum_i p_i(R') \cdot u_i\ \forall\ R' \in \mathcal{R}\}$. That is, $\mathcal{U}^*(p)$ is that subset of $\mathcal{U}$ where the full information allocation under p is utilitarian.

Let $E : \Delta\mathcal{U} \to \mathcal{U}$ denote the expectation operator. Define $\text{Supp}^* : \Delta\Delta\mathcal{U} \rightrightarrows \mathcal{U}$ as $Supp^*(\pi) = E(supp\ \pi)$ for all $\pi \in \Delta\Delta\mathcal{U}$, where $\text{supp} : \Delta\Delta\mathcal{U} \rightrightarrows \Delta\Theta$ is the usual support correspondence. Clearly, for any $\pi \in \Delta\Delta\mathcal{U}$, p is weakly utilitarian over $Supp^*(\pi)$ if and only if $Supp^*(\pi) \subseteq \mathcal{U}^*(p)$.

For all $R \in \mathcal{R}$, define $\mathcal{C}^*_{R,p} = \mathcal{U}(R) \cap \mathcal{U}^*(p)$ and $\mathcal{C}^0_{R,p} = \mathcal{U}(R) \setminus \mathcal{U}^*(p)$.

Lemma A.3.3. *An ordinal rule $p : \mathcal{R} \to \Delta\mathcal{D}$ is weakly utilitarian over $\mathcal{U} = [\underline{u}, \overline{u}]^{m\times n}$ with $0 \leq \underline{u} < \overline{u}$, if and only if p is a constant rule, i.e. $p(R) = p^0$ for some $p^0 \in \Delta\mathcal{D}$. Moreover, if p is not constant, there exists an $R \in \mathcal{R}$ and an open subset $\mathcal{U}^0 \subseteq \mathcal{U}$ such that $\mathcal{U}^0 \subseteq \mathcal{C}^0_{R,p}$.*

Proof. Let $H = \{a_1, \cdots, a_m\}$. Consider the ordinal preference profile P with $a_j P_i a_{j'}$ for all $1 \leq j < j' \leq m, i \in I$, i.e. all agents have the same strict ordinal preference over the objects under P. Fix any rule $p : \mathcal{R} \to \Delta\mathcal{D}$. Let $p(P) = [p_{ij}]_{m\times n}$ where p_{ij} is agent i's probability of getting object a_j under the rule p at preference profile P. Consider any $R \in \mathcal{R}$ such that $p(R) \neq p(P)$. Let $p(R) = [q_{ij}]_{m\times n}$. For any utility profile $u \in \mathcal{U}$, let us define $v_{ij} = (u_{ij} - u_{im})$ for all i, j - agent i's relative preference for object j over object m. Define $w_{ij} = (v_{ij} - v_{nj})$ for all i, j, which captures how much more strongly agent i prefers object a_j to a_m, compared to agent n, at utility profile u.

Using the facts that $\sum_j p_{ij} = 1$ for all i and $\sum_i p_{ij} = q_{a_j}$ for all j, we have:

$$(p(R) - p(P)) \cdot u = \sum_{i\leq n-1, j\leq m-1} (q_{ij} - p_{ij}) w_{ij} \tag{4}$$

We want to construct a $u \in \mathcal{U}$ such that there exists an open ball $B_\epsilon(u)$ around it for some $\epsilon > 0$ so that $(p(R) - p(P)) \cdot u' > 0$ for all $u' \in B_\epsilon(u)$.

Define $I_{-n} = I \setminus n$. Define $= \frac{\overline{u}-\underline{u}}{m+2}$. Let $d = \frac{\Delta}{3}$.

Fix $j \in \{1, \cdots, m-1\}$. Suppose $(q_{ij} - p_{ij}) \geq 0$ for all $i \in I' \subseteq I_{-n}$ and $(q_{ij} - p_{ij}) < 0$ for all $i \in I'' \equiv I_{-n} \setminus I'$. Choose any sequence $\{u_{ij}\}_{i\in I} \subset (\overline{u} - j\Delta - d, \overline{u} - j\Delta + d)$ so that $u_{ij} > u_{nj}$ for all $i \in I'$ and $u_{nj} > u_{ij}$ for all $i \in I''$, $u_{ij} \neq u_{i'j}$ for all $i \neq i'$. Define $d_m = \min_{i\leq n-1, j\leq m-1} |u_{ij} - u_{nj}|$. Choose any sequence $\{u_{im}\}_{i\in I} \subset (\overline{u} - m\Delta - \frac{d_m}{2}, \overline{u} - m\Delta + \frac{d_m}{2})$, $u_{im} \neq u_{i'm}$ for all $i \neq i'$. $\therefore$ $|u_{im} - u_{nm}| < |u_{ij} - u_{nj}|$ for all $i \leq n-1, j \leq m-1$. Fix $i \leq n-1, j \leq m-1$ such that $u_{ij} < u_{nj}$. Therefore $v_{ij} - v_{nj} = (u_{ij} - u_{nj}) - (u_{im} - u_{nm}) > 0$, because in this case $|u_{ij} - u_{nj}| = (u_{ij} - u_{nj}) > 0$, so $(u_{im} - u_{nm}) < (u_{ij} - u_{nj})$. The case for those $i \leq n-1, j \leq m-1$ such that $u_{ij} > u_{nj}$ is analogous.

Clearly, $u \equiv [u_{ij}] \in \mathcal{U}(P)$ is a utility profile for which $(p(R)-p(P))\cdot u > 0$. Define $d^* = \min_{i,j,i',j'} |u_{ij} - u_{i'j'}|$, and $\epsilon = \frac{d^*}{2}$. Therefore for all $u' \in B_\epsilon(u)$, $(p(R)-p(P))\cdot u > 0$. Setting $\mathcal{U}^0 = B_\epsilon(u)$, we have the desired result.

□

Lemma A.3.4. *There exists a finite refinement of π_{ord}, π^+_{ord}, such that p is not weakly utilitarian over $Supp^*(\pi^+_{ord})$.*

Proof. The partition corresponding to π_{ord} is $\{\mathcal{U}(\mathcal{R})\}_{R\in\mathcal{R}}$. Now refine this by splitting each partitional element $\mathcal{U}(R)$ as $\mathcal{C}^*_{R,p} = \mathcal{U}(R) \cap \mathcal{U}^*(p)$ and $\mathcal{C}^0_{R,p} = \mathcal{U}(R) \setminus \mathcal{U}^*(p)$.

By Lemma A.3.3 there exists $R \in \mathcal{R}$ such that $\mathcal{C}^0_{R,p}$ has positive μ-measure. p is not weakly utilitarian with respect to $u^0_R \equiv \mathbb{E}(u|u \in \mathcal{C}^0_{R,p})$ for R, therefore not weakly utilitarian over $Supp^*(\pi^+_{ord})$. □

In order to complete the proof we observe that, by Lemma A.3.4 and Proposition 1, the truthful private ordinal information structure can be the optimal BCE for (p, π^+_{ord}), only if the prior distribution is such that the following holds (recalling the statement of principal's problem (P)):

$$\begin{aligned}
&\mathbb{E}\left(\left(\left(p_i(R_i, R_{-i}) - p_i(R'_i, R_{-i})\right) \cdot \mathbb{E}(u_i|(R_i, R_{-i}))\right)|R_i\right) = 0, \\
\iff &\sum_{R_{-i}\in\mathcal{R}_{-i}} \left((p_i(R_i, R_{-i}) - p_i(R'_i, R_{-i})) \cdot \left(\int\limits_{u_i\in\mathcal{U}_i(R_i)} \frac{u_i d\mu(u_i)}{\mu(\mathcal{U}(R_i, R_{-i}))} \right) \right) \times \frac{\mu(\mathcal{U}(R_i, R_{-i}))}{\sum\limits_{R_{-i}\in\mathcal{R}_{-i}} \mu(\mathcal{U}(R_i, R_{-i}))} = 0 \\
\iff &\sum_{R_{-i}\in\mathcal{R}_{-i}} \left((p_i(R_i, R_{-i}) - p_i(R'_i, R_{-i})) \cdot \left(\int\limits_{u_i\in\mathcal{U}_i(R_i)} u_i \int\limits_{u_{-i}\in\mathcal{U}_{-i}} d\mu(u_i, u_{-i}) \right) \right) = 0
\end{aligned} \tag{5}$$

for some i, R_i, R'_i such that $p_i(R_i, R_{-i}) \neq p_i(R'_i, R_{-i})$ for some R_{-i}.

Equation (5) represents a hyperplane in the space of priors. The (finite) union such hyperplanes across (R_i, R'_i) pairs is a set of measure zero. This establishes part 1 of the statement of Theorem 1.

If, in addition, p is strategyproof, $(p_i(R_i, R_{-i}) - p_i(R'_i, R_{-i})).\mathbb{E}(u_i|(R_i, R_{-i})) > 0$ for all i, R_i, R'_i, R_{-i} such that $p_i(R_i, R_{-i}) \neq p_i(R'_i, R_{-i})$. Hence, $\mathbb{E}\left(\left(\left(p_i(R_i, R_{-i}) - p_i(R'_i, R_{-i})\right).\mathbb{E}(u_i|(R_i, R_{-i}))\right)|R_i\right) > 0$ for all (i, R_i, R'_i) such that $p_i(R_i, R_{-i}) \neq p_i(R'_i, R_{-i})$ for some R_{-i}. By the full support assumption, all $R \in \mathcal{R}$ occur with positive probability. Hence the above equality cannot hold.

Therefore the truthful recommendation signal cannot be the optimal BCE for (p, π^+_{ord}). Therefore by Lemma A.3.2, π_{ord} cannot be the optimal information structure for p.

For the rest of the appendices we use $\mathcal{H} = H^Q$ to denote the multiset containing q_h copies of each object $h \in H$.

A.4 Proofs for Section 4.1

We first define some notation we would need for the proof of Lemma 2 in Section 4.1 as well as for Appendix A.6.

Let $p^* : \mathcal{U} \to \mathcal{D}$ be any deterministic utilitarian allocation rule.[12]

Fix $i \in I, a \in H$ and $u \in \mathcal{U}$ such that an object recommended to agent i by the OR signal with positive probability at u is $a \in H$, i.e. there exists a deterministic utilitarian allocation at $u \in \mathcal{U}$, which allocates a to i.

For all $i \in I, a \in H$, we use $u_{-i,-a}, \mathcal{D}_{-i,-a}$ and $p^*_{-i,-a}$ respectively to denote the restriction of the utility profile $u \in \mathcal{U}$, the set of allowable deterministic allocations $\mathcal{D}$ and the fixed utilitarian allocation rule p^* to $\mathcal{H}' \equiv \mathcal{H} \setminus a, I' \equiv I \setminus i$. Clearly, unless object a has just one copy, $u_{-i,-a}$ is just u_{-i} - the cardinal preferences of agents I' over H.

Proof of Lemma 2. With that background, the following claim is all we need to prove Lemma 2.

Claim A.4.1. *If the posterior top condition is satisfied, reporting his recommended object under the OR signal as his most-preferred object is a best response for each agent, when all other agents obey the recommendation.*

Proof. Fix $i \in I, a \in H$.

Suppose the prior satisfies the posterior top condition. So at any cardinal preference profile u, when agent i has been recommended object a under the OR signal, no matter what he believes the utilitarian allocation of objects in $\mathcal{H} \setminus a$ among agents $I \setminus i$ to be, he weakly prefers his recommended object a to all other objects. Hence, for all $p_i \in \Delta H$ and all $p_{-i,-a} \in \mathcal{D}_{-i,-a}$,

$$\mathbb{E}(u_{ia}|\widehat{p}_i^*(u) = a, p^*_{-i}(u) = p_{-i,-a}) \geq p_i \cdot \mathbb{E}(u_i|\widehat{p}_i^*(u) = a, p^*_{-i}(u) = p_{-i,-a}) \qquad (6)$$

Suppose other agents play some obedient strategy $m^{obed}_{-i}(p_{-i,-a})$ upon observing any joint signal $p_{-i,-a}$, i.e. under $m^{obed}_{-i}(\cdot)$, every agent in $I \setminus i$ reports his recommended object as his most-preferred object.

Consider any strategy $m_i(a)$ played by agent i upon observing recommended object a. The calculations below show that, in expectation, agent i is weakly better off playing any obedient strategy $m^{obed}_i(a)$, i.e. reporting a as his most-preferred object, compared to any other strategy $m_i(a)$.

[12]By our assumption of atomlessness of the prior μ, the utilitarian allocation rule is essentially unique, i.e. two utilitarian welfare-maximizing allocation rules can differ only on a subset of $\mathcal{U}$ of μ-measure zero.

$$\sum_{p_{-i,-a}\in\mathcal{D}_{-i,-a}} \Big[p_i(m_i^{obed}(a), m_{-i}^{obed}(p_{-i,-a}))\cdot \mathbb{E}(u_i|\widehat{p}_i^*(u)=a, p_{-i}^*(u)=p_{-i,-a})\Big]$$
$$\times\, \mu(p_{-i}^*(u)=p_{-i,-a}|\widehat{p}_i^*(u)=a)$$
$$= \sum_{p_{-i,-a}\in\mathcal{D}_{-i,-a}} \mathbb{E}(u_{ia}|\widehat{p}_i^*(u)=a, p_{-i}^*(u)=p_{-i,-a})\mu(p_{-i}^*(u)=p_{-i,-a}|\widehat{p}_i^*(u)=a)$$
$$\geq \sum_{p_{-i,-a}\in\mathcal{D}_{-i,-a}} \Big[p_i(m_i(a), m_{-i}^{obed}(p_{-i,-a}))\cdot \mathbb{E}(u_i|\widehat{p}_i^*(u)=a, p_{-i}^*(u)=p_{-i,-a})\Big]$$
$$\times\, \mu(p_{-i}^*(u)=p_{-i,-a}|\widehat{p}_i^*(u)=a). \text{ (By (6).)}$$

Therefore agent i plays the obedient strategy, i.e. reports the recommended object as his most-preferred object.

□

In order to complete the proof we observe that the mechanism is weakly efficient, and at any utility profile $u \in \mathcal{U}$, any recommended utilitarian allocation is feasible. Therefore when each agent reports his recommended object as his top object, a utilitarian allocation is implemented.

Proof of Theorem 2. By the revelation principle of information design (Bergemann and Morris, 2016), if there exists a signal which induces first-best with a given, fixed mechanism, there exists a direct recommendation signal which induces First Best.[13] Therefore when the mechanism is Deferred Acceptance without Priorities, which is equivalent to Random Serial Dictatorship (RSD), it is without loss to fix the signal space to be the set of all strict ordinal preference ranking profiles $\mathcal{P}$. Let $\nu \in \Delta(\mathcal{P}\times\mathcal{U})$ be the Bayes-correlated equilibrium which implements First Best.

Due to atomlessness of the prior μ, the utilitarian allocation, $p^*(u)$ is unique and deterministic on a subset of $\mathcal{U}$ of μ-measure 1. Therefore the first-best is implementable by RSD implies that the RSD allocation is deterministic with probability one. Clearly, an RSD allocation is deterministic if and only if the set of Pareto efficient deterministic allocations at the reported strict ordinal preference profile is a singleton.

As the next step we use a well-known fact which we also prove for the sake of completeness.

Fact. The set of Pareto efficient deterministic allocations is a singleton at a strict ordinal preference profile only if every agent prefers his allocated object to all other objects at that preference profile.[14]

[13] Bergemann and Morris (2016) consider finite state spaces but it is easy to see that the result holds for state spaces which are compact subsets of metric spaces.

[14] The converse also holds but it is not necessary for our proof.

Proof. Fix $p \in \mathcal{D}$. Suppose $P \in \mathcal{P}$ is such that p is the unique Pareto efficient allocation at P. Fix $i \in I$. Pick any $j \in I \setminus i$ such that $\widehat{p}_i \neq \widehat{p}_j$. Take $p' \in \mathcal{D}$ such that $\widehat{p}_i = \widehat{p}'_j, \widehat{p}_i = \widehat{p}'_j$ and $\widehat{p}_{-i,-j} = \widehat{p}'_{-i,-j}$. By assumption, $\widehat{p}'$ is Pareto inefficient at P. Hence, $\widehat{p}_i P_i \widehat{p}_j$ and $\widehat{p}_j P_j \widehat{p}_i$ (since we assumed strict preferences). By assumption $\sum_h q_h = n$, therefore for all $a \in H \setminus \widehat{p}_i$, there exists $j \in I \setminus i$ such that $\widehat{p}_j = a$. Hence, $\widehat{p}_i P_i a$ for all $a \in H \setminus \widehat{p}_i$, which proves the claim. □

Fix $a \in H$ and $i \in I$ such that $\mu(\widehat{p}^*_i(u) = a) > 0$. Let $\mathcal{P}_i(a)$ be the set of ordinal reports recommended to agent i with positive probability under ν, which have object a on top. Because ν implements first-best, $\nu(P_i \in \mathcal{P}_i(a) | \widehat{p}^*_i(u) = a) = 1$, where P_i denotes i's recommended report. By obedience of reports recommended under ν, when i observes the recommended report P_i, he weakly prefers reporting P_i to reporting any other ordinal preference. But by the above fact, he is indifferent among reporting P_i and any other strict preference with a on top. Therefore all equilibria induced by the OR signal are payoff equivalent to ν, i.e. induce first-best.

Proof of Proposition 2. The fact that $4 \implies 1, 2, 3$ follows directly from Lemma 2. That $2 \implies 4$ follows from Theorem 2. Also obviously, $3 \implies 1, 2$. Below, we only show $1 \implies 4$. This establishes the four-way equivalence.

($1 \implies 4$): Suppose there exists an ordering $\succ$ of agents such that First Best is implementable by Serial Dictatorship by some signal ν. Without loss of generality, let agent 1 be ranked first under the ordering $\succ$. Also without loss of generality, under ν, the message space for agent 1 is the same as his action space, which is the set of objects H. By the atomlessness of μ, the utilitarain welfare-maximizing allocation is unique with probability 1. Because ν implements first-best, at every realized cardinal preference profile, ν recommends to agent 1 the object he should receive under the utilitarain welfare-maximizing allocation, with probability 1. Being the dictator, agent 1's optimal strategy given ν is to pick (one of) his posterior-most-preferred object(s) after receiving any signal. Because ν implements the first-best, at every realized cardinal preference profile, the object agent 1 *should* get under the utilitarian allocation, is (among) agent 1's posterior most preferred object(s). By symmetry, this is true for all other agents as well. Therefore the posterior top condition is satisfied.

A.5 Proofs for Section 4.2

Proof of Proposition 3. *Sufficiency.* Suppose the OR signal is informative on each agent's ordinal preferences. This means, if agent i prefers object b over a a priori, but the realized utility profile u is such that the utilitarian allocation assigns object a to agent i, agent i accepts this recommendation given his information set (which tells him the utilitarian allocation at u). This holds for all i, both objects, and any u. Therefore a posteriori (conditional on receiving the OR signal) each agent always wants the object he is supposed to get under the utilitarian allocation. Therefore First Best is implementable by any weakly efficient mechanism.

Necessity. Suppose there exists i such that i's a priori preference is for object b, and some allocation $p^* \in \mathcal{D}$ with $\widehat{p}_i^* = a$ such that whenever p^* is the publicly recommended allocation under the OR signal - i.e. whenever i knows that p^* is a utilitarian allocation - i continues to prefer b over a conditional on this information. By the full support assumption on the prior, this happens with positive probability. Therefore for such a prior, First Best cannot be implemented whenever agent i is the first dictator under Serial Dictatorship.

Proof of Lemma 3. Let $u \in \mathcal{U}$ be a utility profile such that $p \in \mathcal{D}$ is the unique and deterministic[15] utilitarian allocation at u. Let I_a and I_b be the sets of agents assigned objects a and b respectively under p.

Below, we show that p being the utilitarian allocation at u is equivalent to $rank(\widehat{u}_{i,(a,b)}) \leq q_a$ for all $i \in I_a$ and $rank(\widehat{u}_{i,(b,a)}) \leq q_b$ for all $i \in I_b$.

Necessity. Fix any $i \in I_a, j \in I_b$. Let p' be the allocation obtained by exchanging the allocation of i and j under p and leaving those of the rest of the agents unchanged. By the utilitarianism of p, $p \cdot u \geq p' \cdot u$. This gives $\widehat{u}_{i,(a,b)} \geq \widehat{u}_{j,(a,b)}$. Analogously it follows that $\widehat{u}_{j,(b,a)} \geq \widehat{u}_{i,(b,a)}$ for all $j \in I_b, i \in I_a$. Therefore necessity follows.

Sufficiency. Fix a deterministic allocation p'. Let $I'_a \subseteq I_a$ and $I'_b \subseteq I_b$ be the sets of agents whose allocation changes between p and p'. Clearly $|I'_a| = |I'_b| \leq q_b$. $\widehat{u}_{i,(a,b)} \geq \widehat{u}_{j,(a,b)}$ for all $i \in I'_a, j \in I'_b$. $\therefore$ $\sum\limits_{i \in I'_a} u_{ia} + \sum\limits_{j \in I'_b} u_{jb} \geq \sum\limits_{i \in I'_a} u_{ib} + \sum\limits_{j \in I'_b} u_{ja} \implies$ $\sum\limits_{i \in I_a \setminus I'_a} u_{ia} + \sum\limits_{j \in I_b \setminus I'_b} u_{jb} + \sum\limits_{i \in I'_a} u_{ia} + \sum\limits_{j \in I'_b} u_{jb} \geq \sum\limits_{i \in I_a \setminus I'_a} u_{ia} + \sum\limits_{j \in I_b \setminus I'_b} u_{jb} + \sum\limits_{i \in I'_a} u_{ib} + \sum\limits_{j \in I'_b} u_{ja} \iff$ $p \cdot u \geq p' \cdot u$.

As we saw above, the public OR signal is equivalent to letting each agent know the complete set of agents I_a who have $\widehat{u}_{i,(a,b)}) \leq q_a$. On the contrary, the rank-threshold signal conveys to each agent i only whether $rank(\widehat{u}_{i,(a,b)}) \leq q_a$. However, note that when agent preferences are i.i.d., these two information sets are equivalent. Therefore in that case, $\mathbb{E}(\widehat{u}_{i,(a,b)}|rank(\widehat{u}_{i,(a,b)}) \leq q_a) = \mathbb{E}(\widehat{u}_{i,(a,b)}|\text{public OR signal where the allocation of agent } i \text{ is } a)$ and similarly when the allocation is b. The establishes the equivalence.

Proof of Proposition 4. Suppose, without loss, all agents prefer object a over b a priori. At any utility profile u, all q_b agents allocated b under the utilitarian allocation rule has the same information under the rank threshold signal, therefore the same posterior belief. Therefore, either all of them update their belief to prefer b a posteriori, in which case the First Best allocation for that utility profile is

[15] In other words, u is generic - it lies outside the set of u's of μ-measure zero for which there exist more than one different utilitarian allocations.

implemented, or all of them continue to prefer a a posteriori, in which case the same allocation which would have been implemented under no information ensues. Therefore ex-ante aggregate welfare is weakly higher than in the no information case.

A.6 Proof of Theorem 3

Proof. We proceed in two steps, after which we apply Lemma 2 to conclude the desired result. We use the notation from Appendix A.4, and define the following in addition.

For all $u \in \mathcal{U}$ such that $\widehat{p}_i^*(u) = a$, define $u^*_{-i,-a} = \max_{p \in \mathcal{D}_{-i,-a}} p \cdot u_{-i,-a}$. That is, $u^*_{-i,-a}$ is the maximized social welfare of the economy consisting of agents I' and objects (including copies) $\mathcal{H}'$ at utility profile $u_{-i,-a}$. Clearly, $u^*_{-i,-a} = p^*_{-i,-a}(u) \cdot u_{-i,-a}$, because any utilitarian rule is utilitarian over the subset of agents I' and the set of objects allocated to them under the utilitarian rule, $\mathcal{H}'$. Analogous to agent i's relative preference for object a over b, we can define the rest of the society - I''s - utilitarian relative preference for a over b as $\widehat{u}^*_{-i,-b,-a} \equiv u^*_{-i,-b} - u^*_{-i,-a}$ $\forall\ a, b \in H, i \in I$.

Analogously, for all $i \in I, a, b \in H$ we use $u_{i,-a}$ and $u_{i,-a,-b}$ to denote i's utility profile restricted to objects in $H \setminus a$ and $H \setminus \{a, b\}$ respectively.

Claim A.6.1. *At utility profile $u \in \mathcal{U}$, any deterministic utilitarian welfare-maximizing rule allocates object a to agent i if* [16]*,*

$$u_{i,a} - u_{i,b} > u^*_{-i,-b} - u^*_{-i,-a}\ \forall\ b \in H \tag{UtilSuff}$$

and only if,

$$u_{i,a} - u_{i,b} \geq u^*_{-i,-b} - u^*_{-i,-a}\ \forall\ b \in H \tag{UtilNess}$$

Proof. Fix $u \in \mathcal{U}$ and $i \in I$.

(Necessity) Suppose $\widehat{p}_i^*(u) = a$ for some $a \in H$. Further suppose, by way of contradiction, that (UtilNess) does not hold. Therefore, there exists $b \in H \setminus h$ such that, $u_{i,a} + u^*_{-i,-a} < u_{i,b} + u^*_{-i,-b}$, i.e. the feasible allocation which assigns object b to agent i and chooses a deterministic feasible allocation of objects $\mathcal{H} \setminus a$ among agents $I \setminus i$ which maximizes the total welfare of agents $I \setminus i$, gives strictly greater social welfare than p^* at u. This is a contradiction.

(Sufficiency) Suppose (UtilSuff) holds. We want to show that, $\widehat{p}_i^*(u) = a$. By way of contradiction, suppose there exists $b \in H \setminus a$ such that $\widehat{p}_i^*(u) = b$. Therefore

[16] For our proofs we only need the necessity part but we present the sufficiency conditions as well for the sake of completeness.

by (UtilNess), $u_{i,b} + u^*_{-i,-b} \geq u_{i,a} + u^*_{-i,-a}$. However, by (UtilSuff), $u_{i,a} + u^*_{-i,-a} > u_{i,b} + u^*_{-i,-b}$. This is a contradiction. $\square$

Claim A.6.2. *Under either of conditions (1) and (2), the posterior top condition is satisfied.*

Proof. Fix $p \equiv (p_i, p_{-i}) \in \mathcal{D}$ such that $\widehat{p}_i = a$. Therefore $p_{-i} \in \mathcal{D}_{-i,-a}$. Note that $p^*_{-i}(u) = p_{-i} \Leftrightarrow u_{-i} \in \{u_{-i} \in \mathcal{U}_{-i} : p_{-i} \cdot u_{-i} \geq p'_{-i} \cdot u_{-i} \; \forall \; p'_{-i} \in \mathcal{D}_{-i,-a}\} =: \mathcal{U}_{-i,p_{-i}}(a)$.

Fix any object $b \in H \setminus a$. Below we calculate the interim expected relative preference of agent i for object a over object b, when agent i has been recommended object a, and he knows that the other agents have been jointly recommended the allocation p_{-i}, i.e. when agent i knows u to be such that $p^*(u) = p$. Subsequently we take expectation over all u_{-i} such that $\widehat{p}^*_i(u) = a$ to conclude that $\mathbb{E}\left(\widehat{u}_{i,(a,b)} \mid p^*_i(u) = a\right) \geq 0$.

Fix any $u_{i,-a}, u_{-i}$.

$$
\begin{aligned}
&\mathbb{E}(u_{ia}|p^*(u) = p, u_{i,-a}, u_{-i})\\
&= \mathbb{E}(u_{ia}|p \cdot u \geq p' \cdot u \; \forall \; p' \in \mathcal{D}, u_{i,-a}, u_{-i})\\
&= \mathbb{E}(u_{ia}|u_{ia} \geq \max_{h \in H \setminus a}(u_{ih} + \widehat{u}^*_{-i,-h,-a}), u_{i,-a}, u_{-i}) \text{ (by Claim A.6.1 and atomlessness of } \mu)\\
&\geq \mathbb{E}(u_{ia}|u_{ia} \geq u_{ib} + \widehat{u}^*_{-i,-b,-a}, u_{ib}, u_{i,-a,-b}, u_{-i}) \; (\because \max_{h \in H \setminus a}(u_{ih} + \widehat{u}^*_{-i,-h,-a}) \geq u_{ib} + \widehat{u}^*_{-i,-b,-a})
\end{aligned}
$$

Taking expectation over $u_{i,-a,-b}$ on both sides,

$$
\begin{aligned}
&\mathbb{E}(u_{ia}|p^*(u) = p, u_{ib}, u_{-i})\\
&= \mathbb{E}(u_{ia}|u_{ia} \geq u_{ib} + \widehat{u}^*_{-i,-b,-a}, u_{ia} \geq \max_{h \in H \setminus a,b}(u_{ih} + \widehat{u}^*_{-i,-h,-a}), u_{ib}, u_{-i})\\
&\geq \mathbb{E}(u_{ia}|u_{ia} \geq u_{ib} + \widehat{u}^*_{-i,-b,-a}, u_{ib}, u_{-i}) \qquad (7)
\end{aligned}
$$

Recall equation (Reg) from the definition of Regularity. Putting $H' = H \setminus b$, $\mathcal{I}_{i,H'} = \{u_{i,-b} \in \mathcal{U}_{i,-b} : u_{ia} \geq \max_{h \in H \setminus a,b}(u_{ih} + \widehat{u}^*_{-i,-h,-a}), u_{-i}\}$ and $\mathcal{I} = \{u \in \mathcal{U} : u_{ia} \geq u_{ib} + \widehat{u}^*_{-i,-b,-a}, u_{-i}\}$ in equation (Reg), and using equation (7) we have,

$$
\begin{aligned}
&|\mathbb{E}(u_{ia}|u_{ia} \geq \max_{h \in H \setminus a}(u_{ih} + \widehat{u}^*_{-i,-h,-a}), u_{-i}) - \mathbb{E}(u_{ia}|u_{ia} \geq u_{ib} + \widehat{u}^*_{-i,-b,-a}, u_{-i})|\\
&= \mathbb{E}(u_{ia}|u_{ia} \geq \max_{h \in H \setminus a}(u_{ih} + \widehat{u}^*_{-i,-h,-a}), u_{-i}) - \mathbb{E}(u_{ia}|u_{ia} \geq u_{ib} + \widehat{u}^*_{-i,-b,-a}, u_{-i})\\
&\geq \mathbb{E}(u_{ib}|u_{ia} \geq \max_{h \in H \setminus a}(u_{ih} + \widehat{u}^*_{-i,-h,-a}), u_{-i}) - \mathbb{E}(u_{ib}|u_{ia} \geq u_{ib} + \widehat{u}^*_{-i,-b,-a}, u_{-i})\\
&\Rightarrow \mathbb{E}(u_{ia}|u_{ia} \geq \max_{h \in H \setminus a}(u_{ih} + \widehat{u}^*_{-i,-h,-a}), u_{-i}) - \mathbb{E}(u_{ib}|u_{ia} \geq \max_{h \in H \setminus a}(u_{ih} + \widehat{u}^*_{-i,-h,-a}), u_{-i})\\
&\geq \mathbb{E}(u_{ia}|u_{ia} \geq u_{ib} + \widehat{u}^*_{-i,-b,-a}, u_{-i}) - \mathbb{E}(u_{ib}|u_{ia} \geq u_{ib} + \widehat{u}^*_{-i,-b,-a}, u_{-i})\\
&\Rightarrow \mathbb{E}(\widehat{u}_{i,(a,b)}|\widehat{p}^*_i(u) = a, u_{-i}) \geq \mathbb{E}(\widehat{u}_{i,(a,b)}|\widehat{u}_{i,(a,b)} \geq \widehat{u}^*_{-i,-b,-a}, u_{-i}) \qquad (8)
\end{aligned}
$$

Taking expectation of the above over $u_{-i} \in \mathcal{U}_{-i,p_{-i}}$ we have,

$$\begin{aligned}
&\mathbb{E}(\widehat{u}_{i,(a,b)}|p^*(u) = p) = \mathbb{E}\left(\mathbb{E}(\widehat{u}_{i,(a,b)}|\widehat{p}_i^*(u) = a, u_{-i})|\widehat{p}_i^*(u) = a, p_{-i}^*(u) = p_{-i}\right) \\
&\quad = \mathbb{E}\left(\mathbb{E}(\widehat{u}_{i,(a,b)}|\widehat{p}_i^*(u) = a, u_{-i})|\widehat{p}_i^*(u) = a, u_{-i} \in \mathcal{U}_{-i,p_{-i}}\right) \\
&\quad = \mathbb{E}\left(\mathbb{E}(\widehat{u}_{i,(a,b)}|\widehat{p}_i^*(u) = a, u_{-i})|u_{-i} \in \mathcal{U}_{-i,p_{-i}}\right) \ (\because\ u_i \perp\!\!\!\perp u_{-i}) \\
&\quad \geq \mathbb{E}\left(\mathbb{E}(\widehat{u}_{i,(a,b)}|\widehat{u}_{i,(a,b)} \geq \widehat{u}^*_{-i,(-b,-a)}, u_{-i})|u_{-i} \in \mathcal{U}_{-i,p_{-i}}\right) \ (\text{by } (8))
\end{aligned} \tag{9}$$

Agent i knows u is such that $\widehat{p}_i^*(u) = a$. Fix object $b \in H \setminus a$ and pick any $j \in I$ such that $\widehat{p}_j^*(u) = b$. Therefore $u_{ia} + u_{jb} + p_{-i,-j} \cdot u_{-i,-j} \geq u_{ib} + u_{ja} + p_{-i,-j} \cdot u_{-i,-j}$, i.e. $\widehat{u}_{i,(a,b)} \geq \widehat{u}_{j,(a,b)}$. This is true for all q_b agents $j \in I$ such that $\widehat{p}_j^*(u) = b$. Therefore $\widehat{u}_{i,(a,b)} \geq \max\limits_{\widehat{p}_j^*(u)=b} \widehat{u}_{j,(a,b)}$. Clearly, $\widehat{u}^*_{-i,(-b,-a)} = \max\{\max\limits_{\widehat{p}_j^*(u)=b} \widehat{u}_{j,(a,b)}, \text{ other terms}\} \geq \max\limits_{\widehat{p}_j^*(u)=b} \widehat{u}_{j,(a,b)}$.

Below we show that whenever agent i knows that his utilitarian allocation is object a, a is among his posterior most-preferred objects.

Let $\underline{\widehat{u}}_{r,(a,b)}$ denote the r-th lowest value in $\{\widehat{u}_{j,(a,b)}\}_{j\in I}, r \leq n$.

Taking expectation over u_{-i} such that $\widehat{p}_i^*(u) = a$ in (9),

$$\begin{aligned}
&\mathbb{E}(\widehat{u}_{i,(a,b)}|\widehat{p}_i^*(u) = a) \geq \mathbb{E}\left(\widehat{u}_{i,(a,b)}|\widehat{u}_{i,(a,b)} \geq \widehat{u}^*_{-i,(-b,-a)}\right) \\
&= \mathbb{E}_{u_{-i}}\left(\mathbb{E}_{u_i}\left(\widehat{u}_{i,(a,b)}|\widehat{u}_{i,(a,b)} \geq \widehat{u}^*_{-i,(-b,-a)}, u_{-i}\right)\right) \ (\because\ u_i \perp\!\!\!\perp u_{-i}) \\
&\geq \mathbb{E}_{u_{-i}}\left(\mathbb{E}_{u_i}\left(\widehat{u}_{i,(a,b)}|\widehat{u}_{i,(a,b)} \geq \max_{\widehat{p}_j^*=b} \widehat{u}_{j,(a,b)}, u_{-i}\right)\right) \\
&\quad \left(\because\ \widehat{u}^*_{-i,(-b,-a)} \geq \max_{\widehat{p}_j^*=b} \widehat{u}_{j,(a,b)}, \text{ and } \mathbb{E}_{u_i}\left(\widehat{u}_{i,(a,b)}|\widehat{u}_{i,(a,b)} \geq v, v\right) \text{ is an increasing function of } v\right) \\
&\geq \mathbb{E}_{u_{-i}}\left(\mathbb{E}_{u_i}\left(\widehat{u}_{i,(a,b)}|\widehat{u}_{i,(a,b)} \geq \underline{\widehat{u}}_{q_b,(a,b)}, u_{-i}\right)\right) \ \left(\because\ \max_{\widehat{p}_j^*=b} \widehat{u}_{j,(a,b)} \geq \underline{\widehat{u}}_{q_b,(a,b)}\right) \\
&= \mathbb{E}_{u_{-i}}\left(\mathbb{E}_{u_i}\left(\widehat{u}_{i,(a,b)}|rank(\widehat{u}_{i,(b,a)}) > q_b, u_{-i}\right)\right) \ \left(\because\ \widehat{u}_{i,(a,b)} \geq \underline{\widehat{u}}_{q_b,(a,b)} \Leftrightarrow rank(\widehat{u}_{i,(b,a)}) > q_b\right) \\
&= \mathbb{E}(\widehat{u}_{i,(a,b)}|rank(\widehat{u}_{i,(b,a)}) > q_b) \\
&\geq 0 \ (\text{By no strong a priori preferences})
\end{aligned}$$

The above holds for any $b \in H \setminus a$. Therefore, a is among agent i's posterior most-preferred objects. □

By the above analysis, the posterior top condition is satisfied under each of the two sets of conditions outlined in the theorem. Therefore first best is implementable by any weakly efficient mechanism using the public OR signal, by Lemma 2. □

A.7 Proofs for section 5.3

Proof of Proposition 6. The "if" direction is obvious. For seeing why the other direction holds, fix a mechanism $p : \mathcal{U} \to \Delta\mathcal{D}$. Let $p^*(u)$ denote the set of allocations which maximize the utilitarian welfare at the cardinal preference profile u. Formally, $p^*(u) = \arg\max_{\widetilde{p}\in p(\mathcal{U})} \widetilde{p} \cdot u \subseteq p(\mathcal{U})$. Define the correspondence $\mathcal{U}_p^* : \mathcal{U} \rightrightarrows \mathcal{U}$ as $\mathcal{U}_p^*(u) = p^{-1}(p^*(u)) \subseteq \mathcal{U} \ \forall \ u \in \mathcal{U}$. Pick any prior $\mu \in int\ BIC^p$ such that p is not μ-almost surely weakly utilitarian, i.e. $\mu\left(u \in \mathcal{U}_p^*(u)\right) < 1$.

Clearly, $p^*(u)$ is non-empty, so $\mathcal{U}_p^*(u)$ is also non-empty for all u. Therefore, by the Axiom of Choice we can define a function $T : \mathcal{U} \to \mathcal{U}$ such that $T(u) \in \mathcal{U}_p^*(u)$ for all $u \in \mathcal{U}$.[17]

Now define the signal:

$$\nu_\epsilon(u'|u) = \begin{cases} 1-\epsilon \text{ if } u' = u, \\ \epsilon, \text{ if } u' = T(u), \\ 0, \text{ otherwise,} \end{cases}$$

which is simply an ϵ-mixture of the unconstrained optimal signal $\nu^*(u'|u) = \begin{cases} 1, \text{ if } u' = T(u), \\ 0, \text{ otherwise,} \end{cases}$ with the full information signal $\nu^F(u'|u) = \begin{cases} 1, \text{ if } u' = u, \\ 0, \text{ otherwise.} \end{cases}$

By the interiority of μ, for a sufficiently small ϵ, ν_ϵ is feasible and strictly increases the objective value ($\because \ \mu\left(u \in \mathcal{U}_p^*(u)\right) < 1$), which proves the claim.

A.8 Proofs for Section 6.1

Proof of Proposition 7. Let the two objects be a and b. Without loss, let us assume $q_a \geq q_b$.

Let $\widetilde{p} : \mathcal{R}^D \to \Delta\mathcal{D}$ be a strategyproof and efficient allocation mechanism, where $\mathcal{R}^D$ is the dichotomous domain. Suppose at some preference profile $u \in \mathcal{R}^D$, p is a utilitarian welfare-maximizing deterministic allocation, and $p' \in \ supp(\widetilde{p}(u))$, i.e. p' is the allocation if $\widetilde{p}$ is deterministic, and it is an *ex-post* allocation realized after $\widetilde{p}$ is run on report u, if q involves randomization.

Let I_a and I_b be the sets of agents allocated to objects a and b respectively, under p. Let $I'_a \subseteq I_a$ and $I'_b \subseteq I_b$ be the sets of agents whose allocations are switched between p and p'. Clearly $|I'_a| = |I'_b| \leq q_b$.

By way of contradiction, let us assume p' is not utilitarian welfare-maximizing.

$\therefore p \cdot u > p' \cdot u$, which means $\sum_{i\in I'_a} \widehat{u}_{i,(a,b)} > \sum_{j\in I'_b} \widehat{u}_{j,(a,b)}$. Therefore $\widehat{u}_{i,(a,b)} > \widehat{u}_{j,(a,b)}$ for some $i \in I'_a, j \in I'_b$. In the dichotomous domain this is possible only if $\widehat{u}_{i,(a,b)} = 1$ and $\widehat{u}_{j,(a,b)} = -1$, i.e. i prefers a to b and j prefers b to a at u. $j \in I'_b \subseteq I_b$ (resp. $i \in I'_a \subseteq I_a$), which means j is allocated object a under p' and b under p (resp. i is allocated object b under p' and a under p'). Therefore agent i and j could exchange

[17] We use such a construction of T for simplicity. More generally, we could have defined T as any map $T : \mathcal{U} \to \Delta\mathcal{U}$ such that $T(u) \in \Delta\mathcal{U}_p^*(u)$ for all $u \in \mathcal{U}$, and modified the definition of ν_ϵ accordingly.

the objects allocated to them under p' to make both of them better off. This is a contradiction to the efficiency of $\widetilde{p}$.

Proof of Proposition 8. First note that for any *ordinal* utility space $\mathcal{U}$ each $R \in \mathcal{R}$ corresponds to exactly one $u \in \mathcal{U}$, i.e. there is a one-to-one mapping between $\mathcal{U}$ and $\mathcal{R}$ which captures this, say $T : \mathcal{R} \to \mathcal{U}$. Therefore if $p^* : \mathcal{U} \to \mathcal{D}$ is a deterministic utilitarian mapping, it can be implemented by the ordinal mechanism $p : \mathcal{R} \to \mathcal{D}$ where $p(R) = p^* \circ T(R)$.

The proof of Proposition 8 uses the ordinal mechanism $p \equiv p^* \circ T$ for a given deterministic utilitarian mechanism p^*, and relies on Theorem 1 of Dasgupta and Mishra (2020), which states that every ordinal mechanism which satisfies neutrality and elementary monotonicity (EM) is *ordinally* Bayesian incentive compatible (OBIC), when the prior is independent and uniform over the set of $m!$ possible strict ordinal rankings over objects.[18] As is obvious from the definition of OBIC used by Dasgupta and Mishra (2020), OBIC implies BIC. Below we show that p satisfies both neutrality and EM. While both definitions are available in Dasgupta and Mishra (2020), we define them here once again for the sake of completeness.

A mechanism $p : \mathcal{P} \to \Delta\mathcal{D}$ is *neutral* if for every $p \in \mathcal{P}$ and every permutation $\sigma : H \to H$, $p_{i,h}(P) = p_{i,\sigma(h)}\left(\sigma(P)\right)$ for all i, h, i.e. the "names" of objects do not matter. A utilitarian mechanism obviously satisfies neutrality.

A mechanism p satisfies elementary monotonicity if for every $i \in I$, every $P_{-i} \in \mathcal{P}^{n-1}$, and every $P_i, P_i' \in \mathcal{P}$ such that some objects a and b are ranked consecutively in P_i with aP_ib, and are swapped in P_i' (without changing anything else), we have,

$$p_{ib}\left(P_i', P_{-i}\right) \geq p_{ib}\left(P_i, P_{-i}\right), \text{and } p_{ia}\left(P_i', P_{-i}\right) \leq p_{ia}\left(P_i, P_{-i}\right),$$

i.e. when two consecutively ranked objected are swapped in an agent's report, the agent's shares of those objects change in the "right" directions.

Now we show that in an ordinal rank-value domain, any deterministic utilitarian mechanism satisfies elementary monotonicity. By way of contradiction, suppose not. Note that for a deterministic mechanism p, the only way it can violate EM is, if there exist $i, P_i, P_i', P_{-i}, a, b$ as described above, and $p_{ib} = 1, p_{ib}' = 0$, where we use p_{ih} and p_{ih}' to denote $p_{ih}\left(P_i, P_{-i}\right)$ and $p_{ih}\left(P_i', P_{-i}\right)$ respectively, for all $h \in H$. Suppose this holds. $p_{ib}' = 0$, so there must exist some $c \in H, c \neq b$ such that $p_{ic}' = 1$. By condition (UtilNess) of Claim A.6.1, $p_{ib} = 1 \implies u_{i,b} - u_{i,c} \geq u^*_{I\setminus i,\mathcal{H}\setminus c} - u^*_{I\setminus i,\mathcal{H}\setminus b}$. By the ordinal rank-value property of the domain, for any $c \notin \{a, b\}$, $u_{i,c}' = u_{i,c}$. By the same property, because the rank of b improves from P_i to P_i', $u_{i,b}' > u_{i,b}$. Therefore $u_{i,b}' - u_{i,c}' > u_{i,b} - u_{i,c} \geq u^*_{I\setminus i,\mathcal{H}\setminus c} - u^*_{I\setminus i,\mathcal{H}\setminus b}$ for all c such that bP_ic (Note that P_{-i}, and hence u_{-i}, remains the same in the two reported preference profiles).

[18] In Dasgupta and Mishra (2020)'s setting $n = m$ and $q_h = 1$ for all h. However it is easy to see that nothing changes in their proof even in the more general setting that we use.

$\therefore$ $p'_{ic} = 0$ for all c such that bP_ic, by (UtilNess). Therefore $p'_{id} = 1$ for some d such that dP_ib. Therefore again, by (UtilNess), $u'_{i,d} - u'_{i,b} \geq u'^{*}_{I\backslash i,\mathcal{H}\backslash b} - u'^{*}_{I\backslash i,\mathcal{H}\backslash d} = u^{*}_{I\backslash i,\mathcal{H}\backslash b} - u^{*}_{I\backslash i,\mathcal{H}\backslash d}$. By the ordinal rank-value property, $u'_{id} = u_{id}$. Combined with the fact that $u'_{ib} > u_{ib}$, this implies $u_{i,d} - u_{i,b} > u^{*}_{I\backslash i,\mathcal{H}\backslash b} - u^{*}_{I\backslash i,\mathcal{H}\backslash d}$, i.e. $u_{i,b} - u_{i,d} < u^{*}_{I\backslash i,\mathcal{H}\backslash d} - u^{*}_{I\backslash i,\mathcal{H}\backslash b}$. Therefore p_{ib} cannot be equal to 1 - a contradiction. This shows that any deterministic utilitarian mechanism satisfies EM.

References

A. Abdulkadiroğlu and T. Sönmez. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica*, 66(3):689–701, 1998.

A. Abdulkadiroğlu and T. Sönmez. School choice: A mechanism design approach. *American economic review*, 93(3):729–747, 2003.

D. Almazro, G. Shahatah, L. Albdulkarim, M. Kherees, R. Martinez, and W. Nzoukou. A survey paper on recommender systems. *arXiv preprint arXiv:1006.5278*, 2010.

R. Alonso and O. Câmara. Persuading voters. *American Economic Review*, 106(11): 3590–3605, 2016.

S. Bade. Serial dictatorship: The unique optimal allocation rule when information is endogenous. *Theoretical Economics*, 10(2):385–410, 2015.

D. Bergemann and S. Morris. Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics*, 11(2):487–522, 2016.

D. Bergemann, B. Brooks, and S. Morris. First-price auctions with general information structures: Implications for bidding and revenue. *Econometrica*, 85(1): 107–143, 2017.

S. Bikhchandani. Two-sided matching with incomplete information. Technical report, mimeo, 2014.

A. Bogomolnaia and H. Moulin. A new solution to the random assignment problem. *Journal of Economic theory*, 100(2):295–328, 2001.

A. Bogomolnaia and H. Moulin. Random matching under dichotomous preferences. *Econometrica*, 72(1):257–279, 2004.

A. Bogomolnaia, H. Moulin, and R. Stong. Collective choice under dichotomous preferences. *Journal of Economic Theory*, 122(2):165–184, 2005.

A. Chakraborty, A. Citanna, and M. Ostrovsky. Two-sided matching with interdependent values. *Journal of Economic Theory*, 145(1):85–105, 2010.

J. Chan, S. Gupta, F. Li, and Y. Wang. Pivotal persuasion. *Journal of Economic theory*, 180:178–202, 2019.

Y. Chen and Y. He. Information acquisition and provision in school choice: an experimental study. Technical report, Working paper, 2017.

S. P. Corcoran, J. L. Jennings, S. R. Cohodes, and C. Sattin-Bajaj. Leveling the playing field for high school choice: Results from a field experiment of informational interventions. Technical report, National Bureau of Economic Research, 2018.

M. E. Cortés-Cediel, I. Cantador, and O. Gil. Recommender systems for e-governance in smart cities: State of the art and research opportunities. In *Proceedings of the international workshop on recommender systems for citizens*, pages 1–6, 2017.

S. Dasgupta and D. Mishra. Ordinal bayesian incentive compatibility in random assignment model. *arXiv preprint arXiv:2009.13104*, 2020.

J. C. Ely. Beeps. *American Economic Review*, 107(1):31–53, 2017.

C. R. Featherstone. A rank-based refinement of ordinal efficiency and a new (but familiar) class of ordinal assignment mechanisms. Technical report, Working Paper, 2011.

C. R. Featherstone. Rank efficiency: Investigating a widespread ordinal welfare criterion. Technical report, Technical report, Working Paper, 2014.

D. Gale and L. S. Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.

Gallup. On second thought: Us adults reflect on their education decisions. 2017.

S. Glazerman, I. Nichols-Barrer, J. Valant, and A. Burnett. Presenting school choice information to parents: An evidence-based guide. ncee 2019-4003. *National Center for Education Evaluation and Regional Assistance*, 2018.

A. V. Goldberg and J. D. Hartline. Collusion-resistant mechanisms for single-parameter agents. In *SODA*, volume 5, pages 620–629. Citeseer, 2005.

C. A. Gomez-Uribe and N. Hunt. The netflix recommender system: Algorithms, business value, and innovation. *ACM Transactions on Management Information Systems (TMIS)*, 6(4):1–19, 2015.

J. Grenet, Y. He, and D. Kübler. Decentralizing centralized matching markets: Implications from early offers in university admissions. Technical report, WZB Discussion Paper, 2019.

Y. He, A. Miralles, M. Pycia, and J. Yan. A pseudo-market approach to allocation with priorities. *American Economic Journal: Microeconomics*, 10(3):272–314, 2018.

A. Hylland and R. Zeckhauser. The efficient allocation of individuals to positions. *Journal of Political economy*, 87(2):293–314, 1979.

N. Immorlica, J. WORK, J. LESHNO, I. LO, and B. LUCIER. Information acquisition in matching markets: The role of price discovery. *Available at SSRN*, 2020.

N. Inostroza and A. Pavan. Persuasion in global games with application to stress testing, 2018.

S. Ji and T. Chen. On designing collusion-resistant incentive mechanisms for mobile crowdsensing systems. In *2017 IEEE Trustcom/BigDataSE/ICESS*, pages 162–169. IEEE, 2017.

E. Kamenica and M. Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.

M. A. Khan, K. P. Rath, Y. Sun, and H. Yu. Large games with a bio-social typology. *Journal of Economic Theory*, 148(3):1122–1149, 2013.

A. Klenke. *Probability theory: a comprehensive course*. Springer Science & Business Media, 2013.

E. Lazarova and D. Dimitrov. Paths to stability in two-sided matching under uncertainty. *International Journal of Game Theory*, 46(1):29–49, 2017.

G. Linden, B. Smith, and J. Y. A. Com. Industry report: Amazon. com recommendations: Item-to-item collaborative filtering. In *IEEE Distributed Systems Online*. Citeseer, 2003.

P. Liu. Strategy-proofness of the probabilistic serial rule on sequentially dichotomous domains. 2017.

Q. Liu, G. J. Mailath, A. Postlewaite, and L. Samuelson. Stable matching with incomplete information. *Econometrica*, 82(2):541–587, 2014.

C. Neilson, C. Allende, F. Gallego, et al. Approximating the equilibrium effects of informed school choice. Technical report, 2019.

S. Noda. Strategic experimentation with random serial dictatorship. *Available at SSRN 2956548*, 2021.

P. Penna and C. Ventre. Collusion-resistant mechanisms with verification yielding optimal solutions. In *European Symposium on Algorithms*, pages 708–719. Springer, 2008.

A. E. Roth and M. Sotomayor. Two-sided matching. *Handbook of game theory with economic applications*, 1:485–541, 1992.

R. I. I. Sagaceta. *Worthwhile or Not? Estimating the Impacts of AI-Based Recommendation Systems on Chile's School Choice System*. PhD thesis, Pontificia Universidad Catolica de Chile (Chile), 2020.

P. Shi. Assortment planning in school choice [preliminary draft]. *Fischer vs. University of Texas*, 2016.

I. Taneva. Information design. *American Economic Journal: Microeconomics*, 11 (4):151–85, 2019.

J. Zhang and J. Zhou. Information disclosure in contests: A bayesian persuasion approach. *The Economic Journal*, 126(597):2197–2217, 2016.