

Dynamic assignment without money: Optimality of spot mechanisms

JULIEN COMBE

CREST, Ecole Polytechnique, IP Paris

VLADYSLAV NORA

Department of Economics, Nazarbayev University

OLIVIER TERCIEUX

Paris School of Economics and CNRS

We study a large market model of dynamic matching with no monetary transfers and a continuum of agents who have to be assigned items at each date. When the social planner can only elicit ordinal agents' preferences, we prove that under a mild regularity assumption, incentive compatible and ordinally efficient allocation rules coincide with spot mechanisms. The latter specify “virtual prices” for items at each date and, for each agent, randomly select a budget of virtual money at the beginning of time. When the social planner can elicit cardinal preferences, we prove that under a similar regularity assumption, incentive compatible and Pareto efficient mechanisms coincide with spot menu of random budgets mechanisms. These are similar to spot mechanisms except that, at the beginning of time, each agent chooses within a menu, a distribution over budget of virtual money.

KEYWORDS. Market design, dynamic matching, course allocation.

JEL CLASSIFICATION. C61, C78, D47, D61.

1. INTRODUCTION

In many contexts, agents have to be assigned streams of items when no monetary transfers are allowed. Some examples include the assignment of civil servants—such as teachers—to positions along their career trajectories, the allocation of courses to students from semester to semester, the assignment of spaces in college dorms during university years, the allocation of organs to hospitals waiting for transplants for their sick patients, etc. However, the literature does not provide much guideline on how to design allocation rules in these dynamic contexts.¹ While the class of possible allocation

Julien Combe: julien.combe@polytechnique.edu

Vladyslav Nora: vladyslav.nora@nu.edu.kz

Olivier Tercieux: tercieux@pse.ens.fr

We are grateful to Yeon-Koo Che, Rakesh Vohra, and two anonymous referees for their helpful comments at various stages of this research. Combe is grateful for the support of Labex ECODEC.

¹There have been a number of attempts to define optimal mechanisms in these dynamic contexts. Most of them rely on repeated games structures where preferences are drawn independently and identically distributed (i.i.d.) over time and are separable. This rules out many of the applications we have in mind. See the related literature section below.

rules can potentially be quite large, we show how efficiency and incentive compatibility requirements narrow it down to fairly simple rules that conform well with prevalent practices.

Typically, in the aforementioned situations, a real money market is not allowed, so a “virtual money” market is a natural option.² In practice, agents are often given a budget of virtual money that they can spend at regular intervals of time on items with a high price or can use it to buy cheap items and save money for future use. Hence, the assignment proceeds simply by having a sequence of spot markets.

One example is the course allocation at Columbia Business School (CBS). Until recently, at CBS, lifetime budgets were given upfront and carried over from semester to semester.³ A student could spend her budget equally in each semester, spend most of it on courses in the first semester, or save it for future use.⁴ The prices on courses were set to clear supply and demand for each course. Eventually, the price for a stream of courses simply corresponds to the sum of prices of each course in the stream. Another example is the assignment of teachers to public schools, as is done in France.⁵ Teachers are initially endowed a budget that depends on their characteristics and is used all along their career.⁶ Each year, each teacher can decide to use her budget to transfer to another school, i.e., to “buy” a position in another school. Teachers can use their budget to buy a position in overdemanded schools if they can afford it. For some underdemanded, mainly disadvantaged schools, prices are actually negative, i.e., teachers would receive a bonus if they go to these schools (and stay there for several years). They could then accumulate more tokens to obtain a future assignment at schools that they desire. Here again, the price of a stream of schools along the career trajectory of a teacher is simply the sum of the prices of each school.⁷ Thus, by construction, spot markets have a special linear pricing structure.

One can imagine many other allocation rules. For instance, upon arriving, one could ask an agent her preferences over streams of items and given the reported preferences, allocate the agent a sequence of items from then on. Indeed, in the context

²For studies on static matching problems with virtual money, see, for instance, Hylland and Zeckhauser (1979), Budish (2011), Budish, Cachon, Kessler, and Othman (2017), or He, Miralles, Pycia, and Yan (2018).

³The Wharton School of Business uses a bidding system for courses as well. However, the mechanism used is different: unused budgets from one semester do not carry over to subsequent semesters (see Budish et al. (2017)).

⁴A full description of the allocation process used until recently is given in the “Guide to Bidding” of CBS from 2016. We note though that CBS is now using a different mechanism; see course-match registration (<https://students.business.columbia.edu/records-registration/course-match-registration>).

⁵See Combe, Tercieux, and Terrier (2022) for institutional details on the French teacher assignment scheme.

⁶The initial budget depends on the number of kids, marital situation, and medical condition.

⁷Dynamic assignment schemes with point systems can also be found in other applications. For instance, to incentivize voluntary participation by hospitals in kidney exchange platforms, point systems rewarding hospitals based on their marginal contribution to the platform have been recently adopted by the National Kidney Registry kidney exchange platform (see Agarwal, Ashlagi, Azevedo, Featherstone, and Karaduman (2019)). In addition, the elite French school Ecole Normale Supérieure has been using a point system for the assignment of students to dorms over the years of study.

of course allocation, based on students' (reported) preferences, a university could decide every year to use an allocation rule to assign students to sequences of courses over the full year spanning several semesters. Similarly, teachers who recently graduated could be presented sequences of schools over the following years. With virtual money, one could directly price these streams of items. Since this approach does not impose any linear structure on prices, it may be more permissive than using spot markets, i.e., the allocation rules obtained in this way may not be obtained through spot markets.

We use a large matching market setting with a continuum of agents introduced by [Ashlagi and Shi \(2016\)](#). However, we study a dynamic market where agents are assigned items sequentially, while [Ashlagi and Shi \(2016\)](#) consider static environments. In our framework, agents are present from date 1 through T (the finite horizon), and at each of these dates, they have to be assigned items that perish at the end of the current period. We first consider the case where the mechanism designer can only elicit ordinal preferences over the sequences of items. We show that under a mild regularity assumption, the class of incentive compatible and ordinally efficient allocation rules coincides with the class of spot mechanisms. A spot mechanism works as follows. It specifies virtual prices for items at each date. At the beginning of time, for each agent, it randomly selects a budget of virtual money according to some distribution. Then, at each date, an item is affordable for this agent if her remaining budget is above the virtual price for this item. At this date, the agent is allocated the item of her choice among affordable options. The agent pays the price of the assigned item and the budget is adjusted accordingly. Together with our prior observation that spot mechanisms impose a linear structure on prices, our result shows, perhaps surprisingly, that this linear structure is what is needed when one requires incentive compatibility and ordinal efficiency.

We then consider the case where the mechanism designer can elicit cardinal preferences. Under a similar regularity assumption, we show a corresponding result: the class of incentive compatible and Pareto efficient mechanisms coincides with a class of mechanisms that we call spot menu of random budget (MRB) mechanisms. A spot MRB mechanism is similar to a spot mechanism: it sets prices for each object at each date and will initially draw a budget for each agent. The main difference is that at the beginning of time, each agent is offered a menu of distributions. The distribution chosen in the menu will be used to randomly select an initial budget of virtual money. Then, similarly to spot mechanisms, each agent uses her budget to buy objects at each date.

Our theoretical results provide insights into the types of mechanisms used in practice. As we already underlined, spot mechanisms are used in real-world markets. Of course, since under spot mechanisms, at a given date, agents do not have to express their preferences on what items they are willing to consume at future dates, these mechanisms may be seen as offering simplicity in agents' decision making or accommodating shocks in preferences that may occur in the future. However, given the special structure of pricing underlying these mechanisms, one may wonder about the losses induced by this special structure. Our main result shows that the loss may be small in markets with a fairly large number of agents. Further, while the optimality of spot mechanisms accords

well with their use in practice, it is interesting to note that in some contexts, the dynamic allocation of items is implemented by market mechanisms that differ from spot mechanisms. For instance, as we already mentioned, the Wharton School of Business uses a bidding system for courses where unused budgets from one semester do not carry over to subsequent semesters. We show by means of examples that such mechanisms precluding transfers of budget from one period to the other are inefficient (and, hence, cannot be replicated by spot mechanisms).⁸ More generally, our results shed light on the lack of efficiency of the alternative assignment schemes.

These results also provide a path toward setting up the prices and the budgets in applications where spot markets are in use and where a social planner has a clear objective to optimize. For instance, for the assignment of teachers to public schools in France, one of the main objectives of the administrator/social planner is to ensure that enough experienced teachers are assigned to disadvantaged schools. Maximizing the number of experienced teachers in disadvantaged schools subject to incentive (and efficiency) constraints can then be solved by optimizing over spot mechanisms only. The question then boils down to choices of spot prices for schools and (distribution of) budgets for teachers.

Related literature

Several works have considered market-like mechanisms with token money. The seminal article is Hylland and Zeckhauser (1979), which defines competitive equilibrium with equal income in an environment with fake money. In this context, agents buy probability shares of items, and prices clear the market. Budish (2011) defines a related concept in combinatorial assignment problems such as course allocation. In a continuum model, Che and Kojima (2010) show that the allocation of the random priority mechanism (or random serial dictatorship) can be obtained by setting prices for each object and drawing the budget of fake money of each agent from a uniform distribution that, following Ashlagi and Shi (2016), we name lottery-plus-cutoffs mechanisms.⁹ Importantly, the authors show that random priority is equivalent to the probabilistic serial mechanism of Bogomolnaia and Moulin (2001). Liu and Pycia (2016) and Ashlagi and Shi (2016) prove that the equivalence with random priority extends to large classes of mechanisms. In particular, Ashlagi and Shi (2016) characterize incentive compatible and efficient allocation rules with a continuum of agents when the designer can only elicit ordinal preferences (under the same regularity assumption as ours).¹⁰ They show that the class of incentive compatible and ordinally efficient mechanisms coincides with

⁸While this is a source of inefficiencies, Budish et al. (2017) argue that allowing the transfer of budgets increases decision complexity, since students have to think about how much of their budget they want to reserve for future use.

⁹Che and Kojima (2010) have a “temporal” interpretation of the random priority mechanism to facilitate its comparison with the probabilistic serial mechanism, but it is formally equivalent to our description.

¹⁰Miralles and Pycia (2020) establish a second welfare theorem in assignment problems without transfers.

the class of lottery-plus-cutoffs mechanisms.¹¹ However, all these articles study static settings, whereas we consider a dynamic environment. In particular, we show that the characterization by [Ashlagi and Shi \(2016\)](#) does not extend to our dynamic setup.¹²

There is an extensive literature on dynamic mechanism design problems. Most of the literature focuses on settings in which monetary transfers are allowed (see [Bergemann and Said \(2011\)](#) for a survey). There is a small body of literature on dynamic mechanisms without transfers. [Jackson and Sonnenschein \(2007\)](#) study a general framework for resource allocation in a finite horizon model without discounting in which agents learn all private information at time 0.¹³ They assume that agents' preferences are additively separable and independently distributed across time and agents. The designer's goal is to achieve ex ante Pareto efficient outcomes. To achieve this goal, they build a budget-based mechanism in which each agent announces his preferences and announcements of agents are "budgeted" so that the distribution of preferences announced over the different dates must mirror the underlying distribution of preferences. Hence, the mechanism links the different periods to enforce incentives. Related ideas have been developed and applied to infinite horizon models with discounting where a designer has to repeatedly allocate a single resource to one of multiple agents, whose values are private and i.i.d. across agents and periods (e.g., [Guo, Conitzer, and Reeves \(2009\)](#) and [Santiago, Gurkan, and Sun \(2019\)](#)).¹⁴ The proposed mechanisms share some similarities with our spot mechanisms; in particular, they are based on artificial currencies. For instance, in [Jackson and Sonnenschein \(2007\)](#), each preference ordering is associated with a budget of token money, and announcing a preference ordering has a price that is taken from the associated preference-specific budget.¹⁵ Beyond this type of

¹¹Lottery-plus-cutoffs mechanisms can be implemented using the standard deferred-acceptance mechanism with random priorities. [Shi \(2022\)](#) defines a large class of mechanisms, which includes lottery-plus-cutoffs mechanisms. He provides conditions under which one can implement these mechanisms using either deferred-acceptance, top trading cycle, or serial dictatorship.

¹²Instead, to prove our characterization, we introduce a generalization of their class of lottery-plus-cutoffs mechanisms that we call generalized lottery-plus-cutoffs (GLC) mechanisms. GLC mechanisms also define prices over sequences of items in our case, but draw the budgets according to a general (possibly non-uniform) distribution. Spot mechanisms can be seen as GLC mechanisms where the prices of sequences have a linear structure. We detail the exact connection in Section 4.2.

¹³[Jackson and Sonnenschein \(2007\)](#) are actually more general: they consider a decision problem that is linked with a large number of independent copies of itself. One possible interpretation is that the same problem is repeated a large number of times.

¹⁴These works combine techniques from repeated games ([Abreu, Pearce, and Stacchetti \(1990\)](#), [Fudenberg, Levine, and Maskin \(1994\)](#)) with some of the ideas in [Jackson and Sonnenschein \(2007\)](#) to show how one can approach efficient outcomes when the discount rate is high enough.

¹⁵In some related works, the budget may not be preference-specific and may endow agents with just a single artificial currency budget. For instance, in [Guo, Conitzer, and Reeves \(2009\)](#), agents have a budget of token money. If they have a high valuation for the item today, they can pay the other agent a certain amount of token money to increase their likelihood of obtaining the item today. In turn, the other agent can use the additional tokens later on to increase his likelihood of obtaining the item whenever he will have a high valuation for the item. In a finite horizon model, at the cost of satisfying incentive constraints approximately, [Gorokh, Banerjee, and Iyer \(2017\)](#) offer mechanisms that endow agents with a budget of artificial currency, and organize a static monetary mechanism in each period with payments in the artificial currency.

similarities, our environments differ in important dimensions. The environments these authors consider correspond to a large repetition of independent problems (which is reflected in the assumption that preferences are drawn i.i.d. over time and are separable). This is the cornerstone to ensure that one can link the problems to incentivize agents to report truthfully their preferences when implementing an ex ante efficient allocation. In contrast, our results do not rely at all on any separability or i.i.d. assumptions, and we cannot rely on Jackson and Sonnenschein's (2007) linkage principle. Dropping the separability and i.i.d. assumptions considerably enlarges the set of applications.¹⁶

Our results also relate to the growing literature on dynamic matching. Bloch and Houy (2012) and Kurino (2014) analyze a dynamic version of the housing market with overlapping generations. In their models, the housing side is fixed at the beginning of time and infinitely durable. In dynamic matching infinite horizon stochastic models, Akbarpour, Li, and Gharan (2020), Baccara, Lee, and Yariv (2020), Anderson, Ashlagi, Gamarnik, and Kanoria (2017), and Ashlagi, Burq, Jaillet, and Manshadi (2019) study the trade-off between matching agents immediately or matching them later so as to benefit from market thickening.¹⁷

Last, our analysis is also related to the literature on combinatorial auctions. Indeed, as we already mentioned, spot mechanisms impose a linear structure on prices, and characterize the efficient and incentive compatible mechanisms. For assignment problems with transfers, Kelso and Crawford (1982) show the existence of market-clearing prices (which by definition assume linearity of pricing) provided that agents' preferences satisfy the so-called gross substitutes condition. Hence, under the latter condition, linear pricing allows one to implement efficient allocations (which is generically unique). More generally, Bikhchandani and Mamer (1997) and Bikhchandani and Ostroy (2002) show that for an economy with transfers, for such a result to hold true, duality for the *integer-valued* assignment problem must hold. In particular, with divisible items, the existence of market-clearing prices is ensured. In contrast, first, our result holds without restricting the preferences of the agents and, notably, without imposing any substitute

¹⁶For instance, coming back to our leading examples, students have different sets of choices of courses across semesters, and teachers' preferences on the schools they want to attend today may depend on the school they were assigned to yesterday (for example, because they decided to move near their current school). More generally, preferences over courses or schools in these applications are likely to be persistent across time. Hence, these applications typically violate the assumptions in Jackson and Sonnenschein (2007).

¹⁷More tangentially related to our work, the literature on online resource allocation and online fair division studies the problem of allocating indivisible items arriving over time over a fixed time horizon to a set of agents. The agents' valuations for the item arriving at a given date are known only after the item arrives and are unknown until then. One main question is how the offline setting where items are all available upfront compares with the online setting where items arrive one at a time (e.g., Karp, Vazirani, and Vazirani (1990)). Other works deal with how much envy can be generated in the online context and how it conflicts with efficiency (e.g., Benade, Kazachkov, Procaccia, and Psomas (2018), Zeng and Psomas (2020), and Bogomolnaia, Moulin, and Sandomirskiy (2022)). A difficulty in this literature is how to deal with an uncertain future. One common view is that an adversary selects a distribution of values from which each agent's values are drawn. Results vary depending on the class of distributions that the adversary can select from. In our model, we assume that the distribution of the agents' preferences is known to the designer, and our continuum model rules out uncertainty.

condition. Second, of course, in our setting with a continuum of agents, indivisibilities are ruled out. One may thus naturally wonder if our assumption that there is a continuum of agents buys us our result. As it turns out, in our economy with no transfers, the continuum assumption is not essential—as discussed in Section 7—for our result to hold true. Further, in this section, we provide an example of an economy with a continuum of agents (and violating our regularity assumption), where linear pricing is *with loss of generality*.

Outline

We begin with an example to illustrate the main concepts and results. Then we introduce a benchmark dynamic allocation problem where each agent is assigned a single object in every period. Although this simple model does not capture a variety of the environments described above, it allows for a clear exposition of main ideas. In Section 4, we then proceed to formally define ordinal mechanisms (i.e., mechanisms where agents only report their ordinal preferences) and state our main result in the context of the benchmark model. We also provide the intuition and sketch the proof of the main result. In Section 5, we extend the analysis to cardinal mechanisms. Section 6 introduces the general framework that encompasses our benchmark model and can be applied to many other settings, including, for instance, the allocation of bundles of objects. In particular, it subsumes the dynamic course allocation application discussed in the [Introduction](#). Section 7 concludes with discussions of the model and future research. All proofs are provided in the [Appendix](#).

2. MOTIVATING EXAMPLE

Consider a stylized example of a course allocation problem illustrating our main result. Every semester, a business school offers two courses: mathematics (M) and finance (F). To graduate, a student must complete two semesters, taking one course per semester. We denote a course sequence by a two-tuple (ab) , where $a \in \{M, F\}$ is the course taken in the first semester and $b \in \{M, F\}$ is the course taken in the second semester. For simplicity, we assume that a student can take any combination of the courses over her curriculum. In total, there are four course sequences: (MM) , (MF) , (FM) , and (FF) .¹⁸ We suppose that there is a unit mass of students with arbitrary ordinal preferences over course sequences. We want to assign each student a course sequence depending on her preference. An assignment can be random, meaning that a student can draw a course sequence from a probability distribution. Moreover, it must depend only on the preferences and not the identities of students; thus, any two students with the same preference face the same distribution.

Motivated by the examples in the [Introduction](#), we begin our investigation with the idea of a virtual money market. There are multiple ways to design it in a dynamic environment. For example, one can give each student two separate budgets of artificial

¹⁸For instance, (MM) can be a specialization in mathematics. The general framework introduced in Section 6 allows bundles of courses at each semester and arbitrary constraints on the acceptable course sequences.

TABLE 1. Allocations for separate budgets over semesters

Course Sequence	Allocation for A	Allocation for B
(FF)	0	0
(MF)	3/9	1/9
(FM)	4/9	6/9
(MM)	2/9	2/9

currency, one for every semester. Alternatively, we can give a single budget transferable across semesters. In addition to budgets, another design dimension is the prices for courses. Should there be a separate price for each course every semester? Alternatively, should there be a price for each course sequence? Next, we illustrate how these design choices contribute to the efficiency of the allocation mechanisms and pin down which design works best.

First, consider a mechanism similar to the “course match” mechanism in the Wharton Business School where each student receives, for every semester, a separate budget that is not transferable across semesters. We allow the budgets to be randomly drawn and independent across students and semesters. Suppose that each budget is drawn uniformly from the unit interval. Furthermore, let the (spot) prices of courses be $p_M^1 = 0$, $p_F^1 = 1/3$, $p_M^2 = 0$, and $p_F^2 = 2/3$, where the superscript denotes the semester. When entering the program, each student receives two budget realizations and then optimally uses each budget to buy a course for the corresponding semester. In Table 1, we provide the resulting (ex ante) allocations for student A with ordinal preferences $(MF) \succ (FM) \succ (FF) \succ (MM)$ and for student B with ordinal preferences $(FM) \succ (MF) \succ (FF) \succ (MM)$.

Note that student A obtains (FM) with a positive probability. It happens when her budget at semester 1 is in $[1/3, 1]$ and her budget at semester 2 is in $[0, 2/3)$. Indeed, in that case, student A cannot afford her most preferred sequence (MF) since the price of finance at semester 2 is $p_F^2 = 2/3$, which is above her budget for that semester, but she can afford her second most preferred sequence (FM) . Similarly, student B obtains (MF) when her budget at semester 1 is in $[0, 1/3)$ and her budget at semester 2 is in $[2/3, 1]$. However, if these students were to trade the probabilities of course sequences (FM) and (MF) , the same mass of each course would be allocated in every semester, while the students would improve their allocations.¹⁹ Hence, our first mechanism fails to produce an efficient allocation.

Given that there is no uncertainty over the future preferences or courses, one could argue that our problem is essentially static. Thus, it is natural to treat course sequences as objects and have a single transferable budget drawn when entering the program. Specifically, suppose that each budget is uniform on a unit interval and that instead of assigning a price to each separate course at each semester, we directly price each

¹⁹By “improving,” we mean in a first-order stochastic dominance sense: for each k , students have a weakly higher probability of receiving one of their top k course sequences (and for some k , this probability is strictly higher).

TABLE 2. Allocations with a single budget and no spot prices.

Course Sequence	Allocation 1	Allocation 2
(MM)	0	1/3
(FF)	1/3	2/3
(FM)	1/3	0
(MF)	1/3	0

TABLE 3. Allocations with a single budget and spot prices.

Course Sequence	Probability
(FF)	0
(MF)	1/3
(FM)	1/3
(MM)	1/3

course sequence in the following way: $p_{MF} = 0$, $p_{FM} = 1/3$, $p_{FF} = 2/3$, and $p_{MM} = 1$. Note that unlike in the previous mechanism, these prices cannot be decomposed as the sums of spot prices across semesters.²⁰ At the start of the program, each student receives a single budget realization and optimally spends her budget to buy a course sequence. Allocation 1 in Table 2 specifies an allocation of a student with preferences $(MM) \succ (FF) \succ (FM) \succ (MF)$. This allocation is not efficient. Indeed, the distribution of allocation 2 assigns the same mass of each course in every semester, and the student is better off with this distribution. Hence, this second mechanism is also inefficient.

Finally, consider a spot mechanism combining a single transferable budget and spot prices. Fix a uniform budget distribution and spot prices $p_M^1 = 0$, $p_F^1 = 1/3$, $p_M^2 = 0$, and $p_F^2 = 2/3$. As an example, we derive an allocation of a student with ordinal preferences $(FF) \succ (MF) \succ (FM) \succ (MM)$. If the realized budget is in $[1/3, 2/3)$, then the student will opt for finance in the first semester and then spend her budget on mathematics in the second semester, thus obtaining course sequence (FM) . The probability of such realization is 1/3 and, hence, the probability of (FM) is 1/3. Similarly, we obtain probabilities in Table 3 for each course sequence. It turns out that the allocation rule induced by this mechanism is efficient. Our main result is that such spot mechanisms actually characterize the entire set of incentive compatible and efficient allocation rules in the dynamic environment. In particular, any incentive compatible and efficient allocation rule can be implemented by an appropriate choice of budget distribution and spot prices, and any budget distribution and spot prices induce an incentive compatible and efficient allocation rule.

²⁰There is no vector of spot prices $(p_M^1, p_F^1, p_M^2, p_F^2)$ such that $p_{ab} = p_a^1 + p_b^2$ for each $(ab) \in \{M, F\}^2$. Indeed, if such spot prices existed, we would have $p_{FF} + p_{MM} = p_{FM} + p_{MF}$, which is not true.

3. THE DYNAMIC ALLOCATION PROBLEM

Consider a dynamic version of the allocation problem introduced by Ashlagi and Shi (2016).²¹ There is a continuum of agents, a sequence of T dates, and, at each date t , a finite set of object types O_t . Every date, each agent must be allocated exactly one object, and the set of pure allocations is given by $\mathbf{O} = O_1 \times \cdots \times O_T$. We allow individuals to receive *random allocations*, which are elements of the probability simplex

$$\Delta = \left\{ \mathbf{q} \in \mathbb{R}^{|\mathbf{O}|} : \mathbf{q} \geq 0, \sum_{\mathbf{o} \in \mathbf{O}} q_{\mathbf{o}} = 1 \right\},$$

where $q_{\mathbf{o}} \geq 0$ is the probability of pure allocation $\mathbf{o} \in \mathbf{O}$.

The problem of the social planner is to design a mechanism that allocates objects by taking into account the preferences of agents. We separately study the two types of mechanisms corresponding to the elicited preferences being either ordinal or cardinal. We begin with ordinal mechanisms because all the applications mentioned in the [Introduction](#) involve ordinal preferences and the main argument for the proof in the cardinal case heavily relies on the proof construction in the ordinal case. We extend our results to cardinal preferences in Section 5.²²

4. ORDINAL MECHANISMS

In this section, we assume that the social planner elicits only ordinal preferences over \mathbf{O} . Suppose that the preferences are strict, and let π denote such an ordinal preference, i.e., a permutation of \mathbf{O} , and let Π denote the set of all such preferences. Hence, we allow for arbitrary complementarities in preferences between objects consumed by an agent on different dates. For $h = 1, \dots, |\mathbf{O}|$, we let $\pi(h)$ be the element of \mathbf{O} on the h th place in an agent's ranking according to the preference π . Let F be a commonly known probability distribution over the ordinal preferences, so that $F(\pi)$ is the mass of agents with preferences π . We say that F has *full support* if $F(\pi) > 0$ for every $\pi \in \Pi$.

A social planner allocates objects available at each date among agents, taking into account their reported ordinal preferences. A *mechanism* (or an allocation rule) \mathbf{x} is a mapping from the set of ordinal preferences to a set of random allocations, $\mathbf{x} : \Pi \rightarrow \Delta$. Given mechanism \mathbf{x} , we denote a corresponding random allocation of an agent with

²¹This is a special case of the general model defined in Section 6. This model can accommodate the case where agents are allocated bundles of objects at each date as well as general constraints on the set of available bundles.

²²To implement a random allocation, one must find a corresponding lottery over pure allocations. The Birkhoff–von Neumann theorem states that this is possible in the static one-to-one environment, but when agents are allocated distributions over bundles (or as in our model, over sequences of items), the theorem no longer holds (e.g., Nguyen, Peivandi, and Vohra (2016)). However, in a model with a continuum of agents, this is irrelevant simply because the probability share of getting a certain allocation could be treated as a share of agents getting that allocation.

preference profile π by $\mathbf{x}(\pi) \in \Delta$.^{23,24} We say that a mechanism is *incentive compatible* (IC) if, for any π, π' and each $m = 1, \dots, |\mathbf{O}|$, we have

$$\sum_{k=1}^m x_{\pi(k)}(\pi) \geq \sum_{k=1}^m x_{\pi(k)}(\pi'). \quad (4.1)$$

In other words, a mechanism is incentive compatible if the random allocation obtained by reporting each agent's true preferences first-order stochastically dominates for this agent each random allocation that can be obtained by reporting some other preferences.²⁵ Another requirement that we impose is that it must be impossible for agents to improve their random allocations in the sense of the first-order stochastic dominance by trading their allocation probabilities. Given date t and object $i \in O_t$, let S_{it} be the set of pure allocations with object i at date t , i.e., $S_{it} = \{\mathbf{o} \in \mathbf{O} : o_t = i\}$. We say that a mechanism \mathbf{x} is *ordinally efficient* (OE) if there is no other mechanism \mathbf{x}' such that the following conditions hold:

- (i) For each date t and object type $i \in O_t$, we have

$$\sum_{\pi \in \Pi} \sum_{\mathbf{o} \in S_{it}} x'_{\mathbf{o}}(\pi) F(\pi) = \sum_{\pi \in \Pi} \sum_{\mathbf{o} \in S_{it}} x_{\mathbf{o}}(\pi) F(\pi).$$

- (ii) For each $m = 1, \dots, |\mathbf{O}|$ and for each π , we have $\sum_{h=1}^m x'_{\pi(h)}(\pi) \geq \sum_{h=1}^m x_{\pi(h)}(\pi)$, with a strict inequality for some m and π such that $F(\pi) > 0$.

The first condition requires that at every date, the mass of allocated objects of every type is the same in \mathbf{x} and \mathbf{x}' . The second condition requires that for each agent, the random allocation associated with \mathbf{x}' first-order stochastically dominates for this agent the random allocation associated with \mathbf{x} . We denote the set of all IC and OE mechanisms by $\mathcal{M}_{\text{IC}}^e$.²⁶

Our goal is to characterize the set of IC and OE mechanisms. Incidentally, we will show that these mechanisms are similar to assignment schemes that are used in practice (e.g., for course allocation at universities and for the assignment of teachers to schools in France).

²³Our definition of a mechanism assumes that agents are treated symmetrically, i.e., agents with the same reported ordinal preferences will receive the same random allocation. In particular, the social planner cannot discriminate based on the observed characteristics of agents. However, it is easy to enrich our environment to allow for the observed characteristics of agents. As in [Ashlagi and Shi \(2016\)](#), we would index mechanisms by these observed "types" and focus on mechanisms that treat agents of the same type symmetrically and that are ordinally efficient within types. It is straightforward to extend our results to this richer environment.

²⁴Note that with a continuum of agents and a full support distribution, there is formally no difference between a mechanism and an assignment of random allocations to agents.

²⁵Since the model is ordinal, we use a definition purely based on ordinal preferences. As is well known, this is equivalent to requiring that each agent maximizes his expected utility by reporting his true preferences π for all cardinal representations of π .

²⁶Whereas we do not explicitly introduce object capacities in the model, they appear implicitly in condition (i) of the definition of OE. Indeed, each allocation rule induces a utilization of capacity. Such allocation is OE if the utilized capacities cannot be reassigned in a way that makes agents better off. We provide a detailed discussion of this in Section 7.

4.1 Spot mechanisms and main characterization

The mechanisms used in practice and described in the [Introduction](#) share a common feature: they give a budget of artificial currency to each agent early on and allocate the objects on the spot, i.e., they let agents manage their budget over time to buy some available objects at each date. To capture this feature, we introduce the following definition. Fix a budget distribution with a continuous cumulative distribution function G over $[0, 1]$ and, for each date $t = 1, \dots, T$, prices $p^t = (p_i^t)_{i \in O_t}$ with $p_i^t \geq 0$ for each object $i \in O_t$ available at this date.²⁷ A mechanism \mathbf{x} is a *spot mechanism* if it can be obtained when each agent a makes dynamically optimal choices in the following procedure.

- **Date 1.** Agent a independently draws a budget according to distribution G . Let b_a^1 be the realized budget of agent a . Then a picks an object among the affordable ones, i.e., in $\{i \in O_1 : p_i^1 \leq b_a^1\}$. If a chooses object $i \in O_1$, the budget is adjusted to $b_a^2 := b_a^1 - p_i^1$.
- **Date $t \geq 2$.** Agent a picks an object among the affordable ones, i.e., in $\{i \in O_t : p_i^t \leq b_a^t\}$. If a chooses object $i \in O_t$, the budget is adjusted to $b_a^{t+1} := b_a^t - p_i^t$.

We make two assumptions to guarantee that for each budget realization, the procedure is well defined, inducing a pure allocation of objects: (i) the object prices and the budget distribution are such that there is an affordable pure allocation for each budget realization, i.e., $\min_{\mathbf{o} \in \mathbf{O}} \sum_{t=1, \dots, T} p_{o_t}^t \leq \inf\{z : G(z) > 0\}$; (ii) each agent must choose an object at each date, i.e., the choices where an agent remains unassigned at some dates are not feasible. Given the previous points, dynamic optimality implies that the sequence of choices of agent a corresponds to his most preferred vector $\mathbf{o} = (o_t)_{t=1, \dots, T}$ in \mathbf{O} such that $\sum_{t=1, \dots, T} p_{o_t}^t \leq b_a^1$. Note that, given our assumption of strict preferences, for each agent a , there is a unique such \mathbf{o} .²⁸ Integrating over all possible realizations of budgets given the distribution G , we obtain a corresponding allocation rule \mathbf{x} . We let \mathcal{G}_{sm} denote the set of spot mechanisms.

Note that the definition captures, in particular, the course allocation procedure used at CBS except for the fact that we have not allowed situations where bundles of objects are allocated at each date t . Section 6 presents an extension of our model that captures this aspect as well. It also resembles the procedure of assigning teachers to schools in France described in the [Introduction](#). For an illustration of spot mechanisms, we refer the reader to Section 2 with our motivating example (our last mechanism).

The main result of this section is that spot mechanisms characterize the entire set of incentive compatible and ordinally efficient allocation rules in dynamic environments.

THEOREM 1. *Suppose that the distribution F has full support. A mechanism \mathbf{x} is incentive compatible and ordinally efficient if and only if it is a spot mechanism, i.e., $\mathcal{M}_{\text{IC}}^e = \mathcal{G}_{\text{sm}}$.*

²⁷The requirement that the distribution has a continuous cumulative distribution function is only needed for the results presented in Section 7 when we relax the full support assumption.

²⁸If the agent were to choose sequentially, then a simple backward induction argument together with the strict preferences assumption would also lead to choosing the same unique allocation.

In Section 6, we introduce a general framework with bundles that subsumes our current model. There we also present Theorem 3, which subsumes Theorem 1, the proof of which is provided in Appendix B.

We conclude this subsection with a few comments on our main result. As we already underlined, spot mechanisms are used in real-world markets. However, one can imagine other mechanisms, and, indeed, other types of mechanisms are used in practice. Our result shows that with a continuum of agents, the restriction to spot mechanisms is without loss as long as one wants to achieve ordinally efficient and incentive compatible allocations. Methodologically, this brings some simplification to a designer's problem having a social objective to optimize. Indeed, if the objective is ordinally efficient, then one has to optimize over spot mechanisms, and the question then boils down to the choices of spot prices for items and the distribution of budgets for agents. In addition, our results shed some light on the lack of efficiency of alternative assignment schemes, some of which are used in practice. Indeed, our motivating example in Section 2 illustrates two natural modifications of the spot mechanisms that turn out to be inefficient, i.e., one where separate budgets are drawn independently for each date and one where prices of pure allocations cannot be decomposed into spot prices. In particular, the observation that the latter mechanisms are inefficient turns out to be a core element of the proof of Theorem 1. The following section presents a sketch of this proof.

4.2 Sketch of the proof

Spot mechanisms are a special case of a larger class of mechanisms. Fix a collection of cutoffs $\alpha := (\alpha_{\mathbf{o}})_{\mathbf{o} \in \mathbf{O}} \in [0, 1]^{|\mathbf{O}|}$ and a distribution G over $[0, 1]$. An allocation rule \mathbf{x} is a *generalized lottery-plus-cutoff* (GLC) *mechanism* with parameters $L := (\alpha, G)$ if for every π and $h = 1, \dots, |\mathbf{O}|$, we have

$$x_{\pi(h)}(\pi) = \Pr\left(b < \min_{m=1, \dots, h-1} \alpha_{\pi(m)}\right) - \Pr\left(b < \min_{m=1, \dots, h} \alpha_{\pi(m)}\right),$$

where b is the random budget drawn according to G .²⁹ Plainly, under a generalized lottery-plus-cutoffs allocation rule, each agent a independently draws a budget b_a from distribution G on the unit interval and chooses her favorite pure allocation \mathbf{o} among those with cutoffs below her budget, i.e., in the set $\{\mathbf{o} \in \mathbf{O} : \alpha_{\mathbf{o}} \leq b_a\}$. We denote a GLC mechanism with parameters L by \mathbf{x}^L and denote the set of allocation rules that are GLC mechanisms by \mathcal{G} .

Spot mechanisms are a subclass of GLC mechanisms with a special “linear” structure of cutoffs. Formally, a spot mechanism is a GLC mechanism with parameters $L = (\alpha, G)$ such that there exists a sequence of profiles of nonnegative prices $\mathbf{p} = (p^t)_{t=1, \dots, T}$, where $p^t = (p_i^t)_{i \in O_t}$ for each $t = 1, \dots, T$, satisfying

$$\alpha_{\mathbf{o}} = \sum_{t=1}^T p_{o_t}^t$$

²⁹With a convention that $\min_{m=1, \dots, h-1} \alpha_{\pi(m)} = 1$ if $h = 1$.

for each $\mathbf{o} = (o_1, \dots, o_T) \in \mathbf{O}$. We will say that cutoffs satisfying the above condition are *linear*. When the cutoffs α in the definition of a GLC mechanism are not linear, unlike spot mechanisms, the resulting GLC mechanism cannot be reproduced by allocating objects on the spot. Thus, one can implement a larger set of allocation rules using GLC mechanisms. However, as shown in our motivating example in Section 2, GLC mechanisms need not be ordinally efficient.

In a static environment, i.e., when $T = 1$, Ashlagi and Shi (2016) characterized ordinally efficient and incentive compatible allocation rules as lottery-plus-cutoff mechanisms. Formally, using our above terminologies, an allocation rule \mathbf{x} is a lottery-plus-cutoffs mechanism if it is a GLC mechanism with parameters $L = (\alpha, G)$, where $G = U_{[0,1]}$. Under a lottery-plus-cutoffs allocation rule, each agent a independently draws a budget b_a from the uniform distribution on the unit interval and chooses her favorite pure allocation \mathbf{o} among those in $\{\mathbf{o} \in \mathbf{O} : \alpha_{\mathbf{o}} \leq b_a\}$. Let \mathcal{G}_{AS} be the set of lottery-plus-cutoffs mechanisms. Ashlagi and Shi’s (2016) (AS) characterization result in the static case states that whenever the distribution F has full support, an allocation rule \mathbf{x} is ordinally efficient and incentive compatible if and only if it is a lottery-plus-cutoff mechanism.

As long as we are in a nontrivial dynamic environment, i.e., when $T \geq 2$, their result fails, as made clear in the motivating example presented in Section 2. However, when $T \geq 2$, we can still interpret an allocation $\mathbf{o} \in \mathbf{O}$ as an item in a static environment and use Ashlagi and Shi (2016)’s “static” notion of ordinal efficiency on these items. Call this notion the *AS ordinal efficiency*. This AS ordinal efficiency in our dynamic setting is not natural. It imposes that there is no alternative allocation rule \mathbf{x}' satisfying, for each $\mathbf{o} \in \mathbf{O}$,

$$\sum_{\pi \in \Pi} x'_{\mathbf{o}}(\pi)F(\pi) = \sum_{\pi \in \Pi} x_{\mathbf{o}}(\pi)F(\pi)$$

(where we recall that $\mathbf{o} = (o_t)_t$ is a pure allocation, one for each date) together with condition (ii) in our definition of ordinal efficiency. Typically, reallocation of objects within a period is not allowed.³⁰ For instance, in our motivating example in Section 2, when discussing the second mechanism, \mathbf{x}' (i.e., allocation 2) violates the above condition, while it uses the same mass of each object at each date. Of course, AS ordinal efficiency is weaker than OE, as stated in the following lemma.

LEMMA 1. *If an allocation rule \mathbf{x} is ordinally efficient, then it is AS-ordinally efficient.*

PROOF. If \mathbf{x} is not AS-ordinally efficient, then one can find another allocation \mathbf{x}' such that Condition (ii) of ordinal efficiency is satisfied, and for each $\mathbf{o} \in \mathbf{O}$,

$$\sum_{\pi \in \Pi} x'_{\mathbf{o}}(\pi)F(\pi) = \sum_{\pi \in \Pi} x_{\mathbf{o}}(\pi)F(\pi).$$

Fix an object $i \in O_t$. Clearly, summing the above equalities over all $\mathbf{o} \in S_{it}$ gives us Condition (i) in the definition of ordinal efficiency. Thus, we conclude that \mathbf{x} is not ordinally efficient. □

³⁰Of course, as mentioned, in a static environment where $T = 1$, both notions coincide.

Equipped with Lemma 1, we obtain that in our dynamic environment with $T \geq 2$, one direction of the characterization by Ashlagi and Shi (2016) holds.

PROPOSITION 1. *Suppose that the distribution F has full support. An allocation rule is incentive compatible and ordinally efficient only if it is a lottery-plus-cutoffs mechanism. Formally, $\mathcal{M}_{IC}^e \subset \mathcal{G}_{AS}$.*

As will be explained below, our main result (Theorem 1) can be proved using Proposition 1 together with the following result.

PROPOSITION 2. *Suppose that the distribution F has full support. Fix an ordinally efficient lottery-plus-cutoffs mechanism \mathbf{x}^L with $L = (\boldsymbol{\alpha}, U_{[0,1]})$. Then there exists a linear collection of cutoffs $\bar{\boldsymbol{\alpha}}$ that has the same strict order as $\boldsymbol{\alpha}$, i.e., $(\alpha_{\mathbf{o}} < \alpha_{\mathbf{o}'}) \Rightarrow (\bar{\alpha}_{\mathbf{o}} < \bar{\alpha}_{\mathbf{o}'})$.*

The cornerstone of the proof of Proposition 2 is the following result from the theory of linear inequalities.³¹

LEMMA 2 (Carver 1921). *For an arbitrary matrix \mathbf{A} , $\mathbf{Ax} < \mathbf{0}$ is feasible if and only if $\mathbf{y} = \mathbf{0}$ is the only solution for $\mathbf{y} \geq \mathbf{0}$ and $\mathbf{A}^T \mathbf{y} = \mathbf{0}$.*

To understand how we apply Lemma 2, consider Example 1 below.

EXAMPLE 1. There are two dates and two objects. Consider an allocation rule \mathbf{x} induced by a lottery-plus-cutoffs mechanism with strict cutoffs $(\alpha_{11}, \alpha_{12}, \alpha_{21}, \alpha_{22})$ such that $\alpha_{12} < \alpha_{21} < \alpha_{22} < \alpha_{11} < 1$. Note that these cutoffs are not linear. Indeed, if they were, it would imply that $\alpha_{11} + \alpha_{22} = \alpha_{12} + \alpha_{21}$, which is not possible given the above ordering of cutoffs. Further, these cutoffs are nonlinear in a stronger sense. In the sequel, say that $(\alpha_{11}, \alpha_{12}, \alpha_{21}, \alpha_{22})$ has a strict linear order if there is a linear collection of cutoffs with the same strict ordering. One can show that $(\alpha_{11}, \alpha_{12}, \alpha_{21}, \alpha_{22})$ does not have a strict linear order.³² Below, we use Lemma 2 to show how the lack of strict linear order implies that \mathbf{x} is not ordinally efficient.

We begin by arguing that the existence of a strict linear order for the vector of cutoffs $(\alpha_{11}, \alpha_{12}, \alpha_{21}, \alpha_{22})$ in our example is equivalent to a certain system of linear inequalities being feasible. First, if vector $(\alpha_{11}, \alpha_{12}, \alpha_{21}, \alpha_{22})$ has a strict linear order, then by definition, there is a vector of nonnegative prices $\mathbf{p} = (p_1^1, p_2^1, p_1^2, p_2^2)^T$ such that $\alpha_{ij} > \alpha_{i'j'}$ implies $p_i^1 + p_j^2 > p_{i'}^1 + p_{j'}^2$. Hence, if our cutoffs have a strict linear order, then the following system of strict inequalities is feasible:

$$\begin{aligned} p_1^1 + p_1^2 &> p_2^1 + p_2^2 \\ p_2^1 + p_2^2 &> p_2^1 + p_1^2 \\ p_2^1 + p_1^2 &> p_1^1 + p_2^2. \end{aligned}$$

³¹See Chapter 7 of Schrijver (1986).

³²To see that there is no linear collection of cutoffs $\bar{\boldsymbol{\alpha}}$ with the same strict ordering as $(\alpha_{11}, \alpha_{12}, \alpha_{21}, \alpha_{22})$, proceed by contradiction and assume there is such a collection $\bar{\boldsymbol{\alpha}}$. Then, denoting \mathbf{p} for the associated sequence of profiles of prices, we would have that (a) $\bar{\alpha}_{12} < \bar{\alpha}_{22}$ implies $p_1^1 < p_2^1$ while (b) $\bar{\alpha}_{21} < \bar{\alpha}_{11}$ would imply $p_2^1 < p_1^1$, a contradiction.

We can rewrite the above system in matrix form as $\mathbf{A}\mathbf{p} < 0$, where

$$\mathbf{A} = \begin{pmatrix} -1 & 1 & -1 & 1 \\ 0 & 0 & 1 & -1 \\ 1 & -1 & -1 & 1 \end{pmatrix}.$$

Second, if our cutoffs $(\alpha_{11}, \alpha_{12}, \alpha_{21}, \alpha_{22})$ have no strict linear order, then we show that $\mathbf{A}\mathbf{p} < 0$ is not feasible. Indeed, if there is a vector \mathbf{p} such that $\mathbf{A}\mathbf{p} < 0$, then one could define new linear cutoffs $\tilde{\alpha}$ by setting $\tilde{\alpha}_{ij} = p_i^1 + p_j^2$ for all i and j .³³ By construction of \mathbf{A} , the new cutoffs are in the same order as $(\alpha_{11}, \alpha_{12}, \alpha_{21}, \alpha_{22})$, which is a contradiction to the assumption that cutoffs have no strict linear order. Therefore, the existence of a strict linear order for the cutoffs is indeed equivalent to the feasibility of $\mathbf{A}\mathbf{p} < 0$.

Given that our cutoffs have no strict linear order, $\mathbf{A}\mathbf{p} < 0$ is not feasible. Hence, Lemma 2 guarantees that there exists $\mathbf{y} \geq 0$, $\mathbf{y} \neq 0$ such that $\mathbf{A}^T\mathbf{y} = 0$. In particular, for any $\varepsilon > 0$, $\mathbf{y} = (\varepsilon, 2\varepsilon, \varepsilon)$ is such a solution of $\mathbf{A}^T\mathbf{y} = 0$. It turns out that we can use \mathbf{y} to specify a sequence of bilateral mass transfers that can improve upon a random allocation \mathbf{q} for some agents while keeping the mass of allocated objects of every type constant in every date. In particular, let $y_1 = \varepsilon$ be the probability mass to be transferred from (22) to (11), let $y_2 = 2\varepsilon$ be the probability mass to be transferred from (21) to (22), and let $y_3 = \varepsilon$ be the probability mass to be transferred from (12) to (21). Then $\mathbf{A}^T\mathbf{y} = 0$ implies that if we were to start at any random allocation and could implement these three transfers, the mass of each object at each date must remain the same. For example, consider object 1 at date 1. When we transfer ε from (22) to (11), the mass of the object increases by ε . Its mass does not change when we transfer 2ε from (21) to (22), and its mass decreases by ε when we transfer ε from (12) to (21). So, in total, its mass has not changed after implementing the transfers. Formally, the change of the mass of object i at date t is captured by the negative of the dot product of the corresponding row of \mathbf{A}^T and \mathbf{y} . Now, to show that allocation rule \mathbf{x} is not ordinally efficient, consider an agent whose ordinal preferences are the same as the order of cutoffs, i.e., $(12) \prec (21) \prec (22) \prec (11)$. By the full support assumption, there is a positive mass of such agents. Because the cutoffs are strict, such an agent is assigned a strictly positive probability of each pure allocation. Hence, for sufficiently small $\varepsilon > 0$, we can implement the above sequence of bilateral mass transfers. Moreover, each bilateral transfer moves the probability from a lower to a higher ranked pure allocation according to this agent's preferences (See Figure 1 for an illustration). Hence, after implementing transfers \mathbf{y} , she obtains a dominating (in first-order stochastic dominance) random assignment, while keeping the mass of each object assigned at each date constant. Therefore, the random allocation \mathbf{x} is not ordinally efficient. \diamond

³³Note that the vector \mathbf{p} is not guaranteed to be nonnegative although this is required by our definition of spot mechanism. Moreover, the cutoffs induced by \mathbf{p} may not belong to the unit interval as also required (if these prices are associated with cutoffs of a GLC mechanism). However, as shown in Lemma 4 in the Appendix, a simple normalization of the vector \mathbf{p} (where we add a sufficiently large number to prices to ensure positivity and then multiply prices by a sufficiently small number to ensure that they lie in the unit interval) ensures that both properties hold.




FIGURE 1. Improving mass transfers for an agent with ordinal preferences $(12) < (21) < (22) < (11)$.

The above example illustrates why cutoffs with a strict linear order are needed for a lottery-plus-cutoff mechanism to be ordinally efficient as stated in Proposition 2.³⁴ Then we can use Proposition 1 to deduce that if $\mathbf{x} \in \mathcal{M}_{IC}^e$, then it is induced by a lottery-plus-cutoffs mechanism, i.e., there exists a collection of cutoffs α such that $\mathbf{x} = \mathbf{x}^L$ with $L = (\alpha, U_{[0,1]})$. From Proposition 2, we can deduce that there exists a collection of prices $\mathbf{p} = (p^t)_{t=1, \dots, T}$, where $p^t = (p^t_{i \in O_t})$ for each $t = 1, \dots, T$ and where the collection of linear cutoffs $\bar{\alpha}$ induced by \mathbf{p} has the same strict order as the collection α , i.e., $(\alpha_{\mathbf{o}} < \alpha_{\mathbf{o}'}) \Rightarrow (\bar{\alpha}_{\mathbf{o}} < \bar{\alpha}_{\mathbf{o}'})$.³⁵

However, the GLC mechanism with parameters $(\bar{\alpha}, U_{[0,1]})$ does not generate the same allocation rule as \mathbf{x} . This is because the linear collection of cutoffs $\bar{\alpha}$ has the same strict ordering as α , but need not have the same values. However, using a properly defined distribution G , we can show that the GLC mechanism $L' := (\bar{\alpha}, G)$ is such that $\mathbf{x}^{L'} = \mathbf{x}$ so that the “only if part” of Theorem 1 obtains. While we believe this part of the theorem is surprising, the “if part” of Theorem 1 is a bit more expected, and its proof, which also uses Lemma 2, is relegated to Appendix B.

REMARK 1 (Linear cutoffs and uniform budget distribution). One cannot use a uniform distribution together with linear cutoffs to generate all the incentive compatible and ordinally efficient rules (contrary to the static case studied in Ashlagi and Shi (2016)). To illustrate this, Example 5 in Appendix A provides an ordinally efficient allocation that cannot be implemented by a lottery-plus-cutoffs mechanism (i.e., with a uniform distribution over budgets) with linear cutoffs.

³⁴The argument presented in Example 1 only works with a collection α of strict cutoffs where $(\mathbf{o} \neq \mathbf{o}') \Rightarrow (\alpha_{\mathbf{o}} \neq \alpha_{\mathbf{o}'})$. It is easy to construct examples with an ordinally efficient random allocation that can only be implemented by a lottery-plus-cutoffs mechanism with non-strict cutoffs. In that case, one has to properly build the resulting probability masses to be transferred, and an important part of the proof is devoted to this construction.

³⁵Cutoffs $\bar{\alpha}$ induced by \mathbf{p} means that $\bar{\alpha}_{\mathbf{o}} = \sum_{t=1}^T p^t_{o_t}$ for each $\mathbf{o} = (o_1, \dots, o_T) \in \mathbf{O}$.

5. CARDINAL MECHANISMS

We have studied a dynamic allocation problem where a social planner can only elicit the ordinal preferences of agents. In this section, we extend the analysis to the case where the planner can elicit a complete cardinal preference profile. Our results here are twofold. First, we introduce a new cardinal allocation mechanism tailored to the dynamic environment. Second, we use this mechanism to prove the main result, which resembles the spot market characterization in the ordinal case.

Consider the dynamic allocation problem from Section 3. In contrast to the previous section, here we let agents have cardinal preferences over a set of pure allocations \mathbf{O} represented by utility vector \mathbf{u} , with each coordinate denoting the utility from consuming a corresponding pure allocation. We let U denote the set of all utility vectors inducing strict ordinal preferences and assume that these utility vectors are distributed according to a continuous probability measure F . For a measurable subset $A \subset U$, we let $F(A)$ denote the mass of agents with utility vectors in A .

We follow Ashlagi and Shi (2016) and impose a full relative support assumption on the distribution F . So as to state this condition, let $D := \{\mathbf{u} \in U : \mathbf{u} \cdot \mathbf{1} = 0\}$. One could understand this regularity condition as imposing that, a priori, an agent's relative preference could, with positive probability, take any direction in D .

To formally define our regularity assumption, let us define $\tilde{D} := \{\mathbf{u} \in D : \|\mathbf{u}\| = 1\}$, where $\|\cdot\|$ is the Euclidean norm. Sets U , D , and \tilde{D} are all endowed with standard topologies.³⁶ Let \mathcal{C} be the collection of cones in D .³⁷ We endow \mathcal{C} with the following topology: $\mathcal{C}' \subset \mathcal{C}$ is open if $\mathcal{C}' \cap \tilde{D}$ is open in \tilde{D} . Following Ashlagi and Shi (2016), we say that distribution F has *full relative support* if for any open cone C in \mathcal{C} , $F(\text{Proj}_D^{-1}(C)) > 0$, where $\text{Proj}_D(\cdot)$ stands for the projection of U into D .

REMARK 2. The full relative support assumption is stronger than the full support assumption introduced in the ordinal setting. Again, at an intuitive level, it ensures that F puts positive mass on any direction in D . For instance, this assumption implies that, for a given pure allocation \mathbf{o} , the vector of utilities where agents assign a high (predetermined) level of relative utility to \mathbf{o} must have a positive mass.³⁸ The interpretation here is that each pure allocation \mathbf{o} can, with positive probability, be a “superstar,” i.e., much

³⁶The set $\mathbb{R}^{|\mathbf{O}|}$ is endowed with the topology induced by the Euclidean norm, and D is endowed with the relative topology, i.e., a set is open in D if it is the intersection of an open set in $\mathbb{R}^{|\mathbf{O}|}$ with D . We endow \tilde{D} with the relative topology, i.e., a set is open in \tilde{D} if it is the intersection of an open set in D with \tilde{D} .

³⁷Recall that a cone is a set C such that for all $\lambda > 0$, $x \in C \implies \lambda x \in C$.

³⁸To make this observation precise, let us fix any $c > 1$. Given a vector of utilities \mathbf{u} , let us denote $\mathbf{u}' := (u_{\mathbf{o}} - \frac{\sum_{\mathbf{o}} u_{\mathbf{o}}}{|\mathbf{O}|})_{\mathbf{o}}$ as the normalized vector of utilities. Note that given the values of utilities for $|\mathbf{O}| - 1$ items, the last value is pinned down by the normalization. Without loss of generality, let us assume that this last item is item \mathbf{O} . The full relative support assumption implies that the measure of the set of utilities \mathbf{u} for which, once normalized, $|u'_{\mathbf{o}}| > c|u'_{\tilde{\mathbf{o}}}|$ for all $\tilde{\mathbf{o}} \neq \mathbf{o}$, $|\mathbf{O}|$ is positive. To see why this is true, consider the cone $C_{\mathbf{o}} = \{\mathbf{u} \in D : |u_{\mathbf{o}}| > c|u_{\tilde{\mathbf{o}}}| \text{ for all } \tilde{\mathbf{o}} \neq \mathbf{o}, |\mathbf{O}|\}$. Cone $C_{\mathbf{o}}$ is an open cone. Indeed, $\mathcal{O} = \{\mathbf{u} \in \mathbb{R}^{|\mathbf{O}|} : |u_{\mathbf{o}}| > c|u_{\tilde{\mathbf{o}}}| \text{ for all } \tilde{\mathbf{o}} \neq \mathbf{o}, |\mathbf{O}|\}$ is open in $\mathbb{R}^{|\mathbf{O}|}$. Hence, $\mathcal{O} \cap D \cap \tilde{D} = C_{\mathbf{o}} \cap \tilde{D}$ is open in \tilde{D} , which implies that $C_{\mathbf{o}}$ is an open cone in \mathcal{C} . Thus, $F(\text{Proj}_D^{-1}(C_{\mathbf{o}})) > 0$. This corresponds exactly to the set of utilities \mathbf{u} for which, once normalized, $|u'_{\mathbf{o}}| > c|u'_{\tilde{\mathbf{o}}}|$ for all $\tilde{\mathbf{o}} \neq \mathbf{o}$, $|\mathbf{O}|$.

better than any other pure allocation. This assumption is satisfied in standard multinomial discrete choice models where the distribution of utilities has unbounded support (e.g., the standard logit, the mixed logit, and the probit models), but it may be violated when the distribution over utilities has a bounded support (e.g., Lee (2016) or Che and Tercieux (2019)).

An allocation rule \mathbf{x} is a mapping from utility vectors to random allocations, $\mathbf{x} : U \rightarrow \Delta$. An allocation rule \mathbf{x} is *incentive compatible* if for each $\mathbf{u} \in U$, reporting the true preferences maximizes the expected utility:

$$\mathbf{u} \in \arg \max_{\mathbf{u}' \in U} \mathbf{u} \cdot \mathbf{x}(\mathbf{u}').$$

An allocation rule \mathbf{x} is *Pareto efficient* if there is no other allocation rule \mathbf{x}' such that the following statements hold:

(a) For each date t and object type $i \in O_t$, we have

$$\int_U \sum_{\mathbf{o} \in S_{it}} x'_{\mathbf{o}}(\mathbf{u}) dF = \int_U \sum_{\mathbf{o} \in S_{it}} x_{\mathbf{o}}(\mathbf{u}) dF.$$

(b) For each $\mathbf{u} \in U$, we have $\mathbf{u} \cdot \mathbf{x}'(\mathbf{u}) \geq \mathbf{u} \cdot \mathbf{x}(\mathbf{u})$ and there is a set $A \subset U$ such that $F(A) > 0$ and the inequality is strict for each $\mathbf{u} \in A$

Condition (a) is the analogue of Condition (i) in the definition of ordinal efficiency: the mass of allocated objects at each date remains the same. Condition (b) states that \mathbf{x}' delivers a weakly higher expected utility to every agent and a strictly higher one for a positive mass of agents. In what follows, we introduce a new cardinal mechanism that can be decentralized through a sequence of spot markets, and we use it to characterize the set of incentive compatible and Pareto efficient allocation rules.

A GLC mechanism is an ordinal mechanism and so it is not flexible enough to differentiate cardinal preferences: if ordinal preferences of two agents coincide, then they receive the same allocation. Therefore, we modify a GLC mechanism so as to obtain a mechanism that is responsive to the cardinal preferences of agents. Whereas a GLC mechanism has a single distribution from which each agent independently draws a budget of artificial currency, we now allow agents to choose from a menu of such distributions. We begin with a collection of cutoffs $\alpha := (\alpha_{\mathbf{o}})_{\mathbf{o} \in \mathbf{O}} \in [0, 1]^{|\mathbf{O}|}$ and a collection of distributions $\mathcal{G} := (G_j)_{j \in J}$ over $[0, 1]$.³⁹ Then a random allocation can be constructed by drawing from an agent's ex ante favorite distribution and then choosing the agent's most preferred affordable allocation given her budget realization. For an agent with utility vector $\mathbf{u} \in U$, let $\mathbf{x}^G(\mathbf{u})$ be the expected utility-maximizing random allocation induced by budget distribution G , that is,

$$x_{u(h)}^G(\mathbf{u}) = \hat{G}\left(\min_{m=1, \dots, h-1} \alpha_{u(m)}\right) - \hat{G}\left(\min_{m=1, \dots, h} \alpha_{u(m)}\right)$$

³⁹In this section, we allow distributions to be discrete to simplify the exposition.

for each $h = 1, \dots, |\mathbf{O}|$, where $u(h) \in \mathbf{O}$ is an allocation on h th place in a preference ranking according to utility vector \mathbf{u} ; $\hat{G}(z)$ is a probability that a random budget drawn according to a distribution with cumulative distribution function (c.d.f.) G is strictly below z . An allocation rule \mathbf{x} is a *menu of random budgets* (MRB) mechanism with parameters $L := (\alpha, \mathcal{G})$ if, for every utility vector \mathbf{u} , there is distribution $G_{j(\mathbf{u})} \in \mathcal{G}$ such that $\mathbf{x}(\mathbf{u}) = \mathbf{x}^{G_{j(\mathbf{u})}}(\mathbf{u})$ and

$$G_{j(\mathbf{u})} \in \operatorname{argmax}_{G \in \mathcal{G}} \mathbf{x}^G(\mathbf{u}) \cdot \mathbf{u}.$$

Note that agents with identical ordinal but different cardinal preferences can choose different budget distributions and, hence, receive different random allocations.

Similar to the case of a GLC mechanism, we can introduce a spot version of a MRB mechanism. Fix a sequence of profiles of nonnegative prices $\mathbf{p} = (p^t)_{t=1, \dots, T}$, where $p^t = (p_i^t)_{i \in O_t}$ for each $t = 1, \dots, T$ and a collection of distributions $\mathcal{G} := (G_j)_{j \in J}$ over $[0, 1]$. A mechanism \mathbf{x} is a *spot MRB mechanism* if it can be obtained when each agent a makes dynamically optimal choices in the following procedure.

- **Date 1.** Each agent chooses a distribution from collection \mathcal{G} and independently draws a budget from it. Let b_a^1 be the realized budget of each agent a . Each agent must pick an object among the feasible ones, i.e., in $\{i \in O_1 : p_i^1 \leq b_a^1\}$. If a chooses object $i \in O_1$, the budget is adjusted to $b_a^2 := b_a^1 - p_i^1$.
- **Date $t \geq 2$.** Each agent picks an object among the feasible ones, i.e., in $\{i \in O_t : p_i^t \leq b_a^t\}$. If agent a chooses object $i \in O_t$, the budget is adjusted to $b_a^{t+1} := b_a^t - p_i^t$.

As in Section 4, we make two assumptions: (i) the object prices and the budget distributions in \mathcal{G} are such that there is an affordable pure allocation for each budget realization, i.e., $\min_{\mathbf{o} \in \mathbf{O}} \sum_{t=1, \dots, T} p_{o_t}^t \leq \inf\{z : G(z) > 0\}$ for each $G \in \mathcal{G}$; (ii) each agent must choose an object at each date. Under these assumptions, spot mechanisms always induce a random allocation (each agent is assigned an object in each date). Clearly, spot MRB mechanisms constitute MRB mechanisms with linear cutoffs. Formally, $L := (\alpha, \mathcal{G})$ is a spot MRB mechanism if there exists a sequence of nonnegative profiles of prices $\mathbf{p} = (p^t)_{t=1, \dots, T}$, where $p^t = (p_i^t)_{i \in O_t}$ for each t satisfying $\alpha_{\mathbf{o}} = \sum_{t=1}^T p_{o_t}^t$ for each $\mathbf{o} \in \mathbf{O}$.

The possibility of the spot market implementation of a MRB mechanism is in contrast to the standard *competitive equilibrium with equal income* (CEEI) approach adopted in Ashlagi and Shi (2016). An allocation rule \mathbf{x} is a CEEI with prices $\alpha \in]0, \infty[^{|\mathbf{O}|}$ if for any \mathbf{u} , $\mathbf{x}(\mathbf{u}) \in \operatorname{argmax}_{\mathbf{q} \in \Delta} \{\mathbf{u} \cdot \mathbf{q} : \alpha \cdot \mathbf{q} \leq 1\}$.⁴⁰ We assume that at least one price is induced at random allocation (each agent is assigned an object in each date). Thus, given a profile of prices, agents use a budget of one unit of artificial currency to buy probability shares of pure allocations. A CEEI approach does not fit our dynamic framework because each agent must choose the entire dynamic allocation at the very first date. Nevertheless, it turns out that there is a connection between the two mechanisms. Each CEEI can be implemented as a MRB mechanism, as the following static example illustrates.

⁴⁰We refer to the word “prices” in two different ways. The first refers to prices of each object at each date (p_i^t above) when defining a spot MRB mechanism. The second (α_o above) refers to the prices associated to each pure allocation in the definition of a CEEI.




FIGURE 2. Illustration of budget distributions in Example 2.

EXAMPLE 2. Take a static model with $T = 1$, and consider an economy where each agent is endowed with a single unit of artificial currency and there are three objects with prices of probability shares $\hat{\alpha}_1 = 0$, $\hat{\alpha}_2 = 0.5$, and $\hat{\alpha}_3 = 2$. In the CEEI, an agent chooses an allocation in the probability simplex that maximizes her expected utility subject to a budget constraint. We shall construct a MRB mechanism that induces the same allocation rule as the CEEI above. First, let a collection of cutoffs for the MRB mechanism be given by the above prices, which are normalized to lie inside the unit interval by dividing each price by the highest price, i.e., $\alpha_1 = 0$, $\alpha_2 = 0.25$, and $\alpha_3 = 1$. Second, for each random allocation that is a part of the CEEI, we associate a distribution of a random budget. In particular, for such an allocation \mathbf{x} , let the corresponding distribution $G_{\mathbf{x}}$ assign probability x_i to α_i for $i = 1, 2, 3$. For instance, in each panel of Figure 2, the four allocations $(1/2, 0, 1/2)$, $(1, 0, 0)$, $(0, 1, 0)$, and $(0, 2/3, 1/3)$, corresponding to the vertices of the shaded budget set, give rise to four budget distributions, $(1/2, 0, 1/2)$, $(1, 0, 0)$, $(0, 1, 0)$, and $(0, 2/3, 1/3)$, respectively, where the first number in each of the latter 3-tuples is the probability that the budget is equal to α_1 , the second number, that it is equal to α_2 , and the third number, that it is equal to α_3 . Consider the random allocations that can be obtained by an agent who draws a random budget from each of these distributions and optimally chooses his pure allocation given budget realization. We illustrate them by the black dots in Figure 2. For instance, an agent who chooses the random budget distribution $(0, 1, 0)$ gets 0.25 units of artificial currency with probability 1. In that case, he can buy either object 1 or 2 under the resulting MRB. Hence, depending on his preferences, the agent will choose one of these two pure allocations represented by the bottom two black allocations in Figure 2(b). Similarly, an agent who chooses the budget distribution $(1/2, 0, 1/2)$ obtains a null budget with probability 1/2. In that case, he can only buy object 1 under the resulting MRB. With probability 1/2, he

receives a budget of 1 and can buy any of the available objects. The random allocation induced by his optimal choices (integrating over all possible realizations of the budget) correspond to one of the three black dots in Figure 2(c). For instance, if the agent prefers 2 over 1 over 3, his optimal choices will generate random allocation $(1/2, 1/2, 0)$. Notice that the random allocations that can be generated by the choice of a random budget distribution all lie inside the CEEI budget set represented by the gray region in Figure 2. Hence, if an agent receives an allocation in the CEEI, then this agent weakly prefers the random budget distribution generated by this allocation to any distribution generated by another allocation and obtains this allocation in the MRB mechanisms with the above menu of budgets and prices. Hence, this MRB mechanism induces the same allocation rule as the CEEI. \diamond

The following result generalizes the observation in the example.

PROPOSITION 3. *If \mathbf{x} is CEEI, then \mathbf{x} is a MRB mechanism.*

PROOF. Suppose \mathbf{x} is a CEEI with prices $\hat{\alpha} = (\hat{\alpha}_{\mathbf{o}})_{\mathbf{o} \in \mathbf{O}}$. Let

$$\alpha_{\mathbf{o}} := \frac{\hat{\alpha}_{\mathbf{o}}}{\max_{\mathbf{o}}\{\hat{\alpha}_{\mathbf{o}} : \hat{\alpha}_{\mathbf{o}} \neq \infty\} + 1}$$

for each $\mathbf{o} \in \mathbf{O}$ such that $\hat{\alpha}_{\mathbf{o}} \neq \infty$, and $\hat{\alpha}_{\mathbf{o}} = 1$ for each $\mathbf{o} \in \mathbf{O}$ such that $\hat{\alpha}_{\mathbf{o}} = \infty$,⁴¹ and let $\alpha = (\alpha_{\mathbf{o}})_{\mathbf{o} \in \mathbf{O}}$. For each $\mathbf{u} \in U$ and $\mathbf{x}(\mathbf{u}) \in \Delta$, let $G_{\mathbf{x}(\mathbf{u})}$ be a discrete distribution that assigns probability $x_{\mathbf{o}}(\mathbf{u})$ to the budget value $\alpha_{\mathbf{o}}$ for each \mathbf{o} , and let $\mathcal{G} = (G_{\mathbf{x}(\mathbf{u})})_{\mathbf{u} \in U}$ be a collection of such distributions. We show that \mathbf{x} is a MRB mechanism with $L = (\alpha, \mathcal{G})$.

Fix a discrete distribution of a random budget G . By choosing some affordable allocation at each realization of a random budget, we induce some ex ante distribution over allocations. Define a feasible choice rule to be a function that chooses an affordable pure allocation for each realization of a random budget. Formally, a feasible choice rule is a function $\psi : [0, 1] \rightarrow \mathbf{O}$ such that $\alpha_{\psi(z)} \leq z$ for any $z \in [0, 1]$. Then, given distribution G , let the set of random allocations that can be induced by some feasible choice rule be

$$B(G) = \left\{ \mathbf{y} \in \Delta : \text{there exists feasible } \psi \text{ such that } y_{\mathbf{o}} = \sum_{z:\psi(z)=\mathbf{o}} P_G(z) \text{ for each } \mathbf{o} \in \mathbf{O} \right\},$$

where $P_G(z)$ is the probability of realization z given G .

Now, if an agent with utility \mathbf{u} optimally chooses an affordable bundle for each realization of the random budget $G_{\mathbf{x}(\mathbf{u})}$, then by construction of $G_{\mathbf{x}(\mathbf{u})}$, the induced ex ante distribution is $\mathbf{x}(\mathbf{u})$. Hence, $\mathbf{x}(\mathbf{u}) \in B(G_{\mathbf{x}(\mathbf{u})})$. Next, we show that if $\mathbf{y} \in B(G_{\mathbf{x}(\mathbf{u})})$, then random allocation \mathbf{y} also belongs to the original budget set in the CEEI mechanism with the collection of prices $\hat{\alpha}$, i.e., $\sum_{\mathbf{o}} y_{\mathbf{o}} \hat{\alpha}_{\mathbf{o}} \leq 1$. Therefore, when choosing from a collection of random budgets \mathcal{G} , it is optimal for an agent with utility \mathbf{u} to choose distribution $G_{\mathbf{x}(\mathbf{u})}$.

⁴¹Note that a probability share of each pure allocation with infinite price is zero for all agents.

Suppose $\mathbf{y} \in B(G_{\mathbf{x}(\mathbf{u})})$ and let ψ be a feasible choice rule that induces \mathbf{y} . We have

$$\begin{aligned} \sum_{\mathbf{o}} \hat{\alpha}_{\mathbf{o}} y_{\mathbf{o}} &= \sum_{\mathbf{o}} \hat{\alpha}_{\mathbf{o}} \sum_{\mathbf{o}': \psi(\alpha_{\mathbf{o}'}) = \mathbf{o}} x_{\mathbf{o}'}(\mathbf{u}) \\ &= \sum_{\mathbf{o}} \sum_{\mathbf{o}': \psi(\alpha_{\mathbf{o}'}) = \mathbf{o}} \hat{\alpha}_{\psi(\alpha_{\mathbf{o}'})} x_{\mathbf{o}'}(\mathbf{u}). \end{aligned}$$

Note the above sum consists of terms $\hat{\alpha}_{\psi(\alpha_{\mathbf{o}'})} x_{\mathbf{o}'}(\mathbf{u})$, and each term enters the sum only once. Hence, we can rewrite it as the sum of all these terms:

$$\begin{aligned} \sum_{\mathbf{o}} \sum_{\mathbf{o}': \psi(\alpha_{\mathbf{o}'}) = \mathbf{o}} \hat{\alpha}_{\psi(\alpha_{\mathbf{o}'})} x_{\mathbf{o}'}(\mathbf{u}) &= \sum_{\mathbf{o}'} \hat{\alpha}_{\psi(\alpha_{\mathbf{o}'})} x_{\mathbf{o}'}(\mathbf{u}) \\ &\leq \sum_{\mathbf{o}'} \hat{\alpha}_{\mathbf{o}'} x_{\mathbf{o}'}(\mathbf{u}) \\ &\leq 1. \end{aligned}$$

Here, the first inequality follows from ψ being a feasible choice rule, i.e., $\alpha_{\psi(\alpha_{\mathbf{o}'})} \leq \alpha_{\mathbf{o}'}$, and the fact that $\alpha_{\mathbf{o}'}$ is just a scaling of $\hat{\alpha}_{\mathbf{o}'}$. The final inequality follows from $\mathbf{x}(\mathbf{u})$ being CEEI. \square

Hence, if the prices in CEEI are linear, then from the above argument, it follows that CEEI can be decentralized using a spot MRB mechanism. Our main result in this section is a cardinal version of Theorem 1.

THEOREM 2. *Suppose that distribution F is continuous and has full relative support. A mechanism \mathbf{x} is incentive compatible and Pareto efficient if and only if it is a spot MRB mechanism.*

In the next section, we generalize the cardinal model to the environment with bundles. In this general environment, Theorem 2 is subsumed by Theorem 4 whose proof is given in Appendix C.

6. THE GENERAL FRAMEWORK

Throughout the analysis, we have focused on a simple dynamic environment where agents are assigned a single object at every date. Although this model describes applications, such as the assignment of teachers to jobs and students to dormitories, it does not address all the situations where bundles of objects are allocated. For instance, in our motivating example of course allocation, students can typically take some number of electives per semester. Moreover, some courses can be pre- or anti-requisites to other courses, and students may be required to earn a certain number of credits over the years to graduate. So as to capture this as well as a variety of other settings, we generalize our benchmark model and state the two theorems that subsume Theorems 1 and 2.

The general model

Fix a finite set of generalized object types O . Each agent must be allocated a feasible (nonempty) bundle of objects. We denote the set of all feasible bundles by $B \subset 2^O$ and write $i \in b$ to denote that bundle $b \in B$ contains object type $i \in O$. We impose two restrictions on B . First, any two bundles in B must have the same size. Second, each bundle in B contains at most one object of each type. A set of random allocations is

$$\Delta = \left\{ \mathbf{q} \in \mathbb{R}^{|B|} : \mathbf{q} \geq 0, \sum_{b \in B} q_b = 1 \right\}.$$

Ordinal preferences

Agents have ordinal strict preferences over B . As before, π denotes such a preference, and Π is the set of all preferences, while $\pi(h) \in B$ is the bundle in h th place in the ranking according to $\pi \in \Pi$. Let F be a probability distribution over ordinal preferences of agents with full support. As before, $F(\pi)$ will denote the mass of agents with preferences π . An allocation rule \mathbf{x} is a mapping from a set of ordinal preferences to a set of random allocations, i.e., $\mathbf{x} : \Pi \rightarrow \Delta$. Definitions of incentive compatibility and ordinal efficiency are similarly adapted to the bundle framework. An allocation rule \mathbf{x} is *incentive compatible* if for any $\pi, \pi' \in \Pi$ and each $m = 1, \dots, |B|$, we have

$$\sum_{k=1}^m x_{\pi(k)}(\pi) \geq \sum_{k=1}^m x_{\pi(k)}(\pi').$$

An allocation rule \mathbf{x} is *ordinally efficient* if there is no other allocation rule \mathbf{x}' such that the following statements hold:

- (i) For each object type $i \in O$, we have

$$\sum_{\pi \in \Pi} \sum_{b: i \in b} x'_b(\pi) F(\pi) = \sum_{\pi \in \Pi} \sum_{b: i \in b} x_b(\pi) F(\pi).$$

- (ii) For each $m = 1, \dots, |B|$ and each $\pi \in \Pi$, we have $\sum_{h=1}^m x'_{\pi(h)}(\pi) \geq \sum_{h=1}^m x_{\pi(h)}(\pi)$, with a strict inequality for some m and π such that $F(\pi) > 0$.

We denote the set of incentive compatible and ordinally efficient allocation rules by \mathcal{M}_{IC}^e .

The above model encompasses our benchmark dynamic allocation model with ordinal preferences. Recall that the dynamic model begins with a finite set of object types O_t for each date t . Without loss of generality, we can let types O_t be disjoint sets. The set of pure allocations was a product $\mathbf{O} = O_1 \times \dots \times O_T$. Now define the corresponding set of generalized object types to be $O = O_1 \cup \dots \cup O_T$. Moreover, a bundle is feasible if and only if it contains exactly one object from each O_t . Then the set of pure allocations \mathbf{O} corresponds to the set of admissible bundles.

EXAMPLE 3. The generalization allows us to include into our benchmark model the possibility of allocating bundles and arbitrarily restricting feasible allocations. As an example, consider a course allocation problem with two semesters and three courses a , b , and c . Suppose that each course is available in both semesters, but that course a is a prerequisite for course c and the same course cannot be taken twice. Moreover, to graduate, each student is required to take two courses. We can model this situation by letting $O = \{a_1, a_2, b_1, b_2, c_1, c_2\}$, where a subscript denotes a semester at which a course is taken. The corresponding set of feasible bundles is $B = \{(a_1, b_1), (a_1, b_2), (a_2, b_2), (b_1, a_2), (a_1, c_2)\}$. \diamond

As before, our goal is to characterize all incentive compatible and ordinally efficient allocation rules. To do so, we now introduce the appropriately modified version of a GLC mechanism. Fix a collection of cutoffs $\alpha := (\alpha_b)_{b \in B} \in [0, 1]^{|B|}$ and a distribution G over $[0, 1]$. An allocation rule \mathbf{x} is a *generalized lottery-plus-cutoff* (GLC) mechanism with parameters $L := (\alpha, G)$ if for every π and $h = 1, \dots, |B|$,

$$x_{\pi(h)}(\pi) = \Pr\left(b < \min_{m=1, \dots, h-1} \alpha_{\pi(m)}\right) - \Pr\left(b < \min_{m=1, \dots, h} \alpha_{\pi(m)}\right),$$

where b is a random budget drawn according to G . We denote a GLC mechanism with parameters $L = (\alpha, G)$ by \mathbf{x}^L . Cutoffs α are *linear* if there exist object prices $\mathbf{p} = (p_i)_{i \in O} \in \mathbb{R}^{|O|}$ such that

$$\alpha_b = \sum_{i \in b} p_i$$

for each $b \in B$. Let \mathcal{G}_L be the set of all GLC mechanisms with linear cutoffs. Now we are ready to state our main result.

THEOREM 3. *Suppose that the distribution F has full support. An allocation rule is incentive compatible and ordinally efficient if and only if it is a GLC mechanism with linear cutoffs, i.e., $\mathcal{M}_{IC}^e = \mathcal{G}_L$.*

As we have already seen, our dynamic framework is embedded into the current one so that Theorem 1 is a corollary of Theorem 3. The sketch of the proof is similar to the one we presented in Section 4.2. The actual proof is provided in Appendix B.

Cardinal preferences

The generalization for cardinal preferences is the mirror analogue of the previous section. Agents have cardinal preferences over B , and we let \mathbf{u} be the utility vector where each coordinate gives the utility for a bundle in B . The distribution F over cardinal utility vectors can also be easily generalized, and the full relative support definition does not change from the one given in Section 5. An allocation rule \mathbf{x} now maps the set U of cardinal utility vectors to Δ , the set of random allocations. The definitions of incentive compatibility and Pareto efficiency can easily be adapted from Section 5. We can similarly modify the definition of a MRB mechanism with parameters $L := (\alpha, \mathcal{G})$ to fit

in this new framework where the collection of cutoffs α is defined for the bundles in B . Cutoffs α are linear if $\alpha_b = \sum_{i \in b} p_i$ for each $b \in B$ and some vector of nonnegative prices $\mathbf{p} = (p_i^t)_{i \in O}$. We now state the generalization of Theorem 2 to the setting with bundles.

THEOREM 4. *Suppose that the distribution F is continuous and has full relative support. A mechanism \mathbf{x} is incentive compatible and Pareto efficient if and only if it is a MRB mechanism with linear cutoffs.*

The proof is relegated to Appendix C.

7. DISCUSSIONS

Capacity constraints

We do not introduce capacities because in many contexts, including our main application of course allocation, these are endogenous choice variables rather than hard constraints. Indeed, for each course, a university can set a target capacity or a desired maximum enrollment, but can potentially enroll more students (see Budish et al. (2017)). This approach provides a greater flexibility by allowing the designer to set up an optimization problem where capacity utilization is endogenous. As we discuss in the next subsection (Designer's problem), our main result helps to substantially simplify such optimization problems.

A recent literature in school choice also adopts this approach. For example, Ashlagi and Shi (2016) emphasize the cost of public school busing when determining the allocation of students. The budget limit imposes complex constraints on the schools, e.g., with the enrollment of students coming from areas that require busing, capacity constraints are tighter. In this context, the exact capacity constraint is also endogenous to the matching: the administration may be willing to expand enrollment (of students living relatively far from the school) at a financial cost (public busing cost). For daycare assignment, Kamada and Kojima (2023) present a similar "budget constraint": by law, the number of caregivers is higher for younger children than for older ones; thus, a daycare capacity depends on the distribution of ages of the accepted children.

Nonetheless, hard capacity constraints can still be included in our model without changing the main insight that OE and IC mechanisms are spot mechanisms. Consider our generalized model of Section 6. For each generalized object $i \in O$, let c_i be the capacity of object i . We let $\mathbf{c} = (c_i)_i$ be the vector of capacities. We assume that there is a null object $\emptyset \in O$ that has infinite capacity ($c_\emptyset = \infty$), which one can interpret as "staying unassigned." We say that an allocation \mathbf{x} is *feasible* for capacities \mathbf{c} if, for each object type $i \in O$,

$$\sum_{\pi \in \Pi} \sum_{b: i \in b} x_b(\pi) F(\pi) \leq c_i.$$

A feasible allocation \mathbf{x} is *ordinally efficient for capacities \mathbf{c}* (OE_c) if there exists no other feasible allocation \mathbf{x}' such that for each $m = 1, \dots, |B|$ and for each π , we have $\sum_{h=1}^m x'_{\pi(h)}(\pi) \geq \sum_{h=1}^m x_{\pi(h)}(\pi)$, with a strict inequality for some m and π such that

$F(\pi) > 0$.⁴² It is immediate to see that, for a feasible allocation, OE_c implies OE; thus, Theorem 3 implies the following corollary.⁴³

COROLLARY 1. *Assume that F is full support. If x is a feasible IC and OE_c mechanism, then it is a spot mechanism.*

Designer's problem

Above, we argued that in many real-life problems, capacities are part of the choice variables of the designer. Hence, in these contexts, it is natural to think of the designer's problem as an optimization problem where capacity utilization is endogenous. Here we provide an example of the type of optimization/mechanism design problem we have in mind. As we will explain, our main result in this paper may be useful to simplify such optimization problems. While this is illustrated through a specific objective of the designer, it will be clear that the argument applies beyond this specific objective.⁴⁴

For a mechanism \mathbf{x} , define the capacity utilization by $q_{i,t} := \sum_{\pi \in \Pi} \sum_{\mathbf{o} \in S_{it}} x_{\mathbf{o}}(\pi) F(\pi)$. Based on the above discussion, let $C_{i,t}(q_{i,t})$ denote the cost incurred by the designer for a capacity utilization $q_{i,t}$ of item i at date t ; the total cost simply adds up these costs across items and dates. For the sake of the example, consider the problem

$$\min_{\mathbf{x}} \sum_{\pi \in \Pi} \sum_{k=1}^{|\mathcal{B}|} k x_{\pi(k)}(\pi) F(\pi) + \sum_{i,t} C_{i,t}(q_{i,t})$$

subject to

(IC) and (OE).

That is, the designer cares about agents' welfare (measured by the average ranks of agents) and capacity utilization cost. The problem has an exponential number of variables and constraints. Specifically, assuming there are n available objects at each date, for each of the $(n^T)!$ ordinal preference profiles, we need to specify n^T probabilities, hence, $n^T \times (n^T)!$ variables, and $O((n^T)^2)$ complex constraints.

By our main result, the above optimization problem is equivalent to a simpler unconstrained program of optimizing over the spot mechanisms. In this problem, there are only $Tn + n^T$ variables, which correspond to spot prices and budget distribution.⁴⁵ Thus,

⁴²This definition is similar to previous definitions of ordinal efficiency except that condition (i) is now replaced by feasibility.

⁴³Using a technique similar to that in the proof of Theorem 3, one can show that the prices of each good under the spot mechanism can be set to zero for the goods that are underutilized, i.e., those for which total allocated mass is strictly less than their capacity.

⁴⁴For instance, we could have handled objectives similar to those in Ashlagi and Shi (2016), where the social planner maximizes a linear combination of utilitarian welfare and max-min welfare subject to other constraints (such as respecting a target budget associated with capacity utilization/public school busing).

⁴⁵The choice of budget distribution G can be reduced to a choice of n^T mass points. Indeed, sort the n^T cutoffs in ascending order. Then each G that puts the same probability mass in between each consecutive cutoff induces the same allocation.

our characterization significantly reduces the number of variables and constraints. While the computational analysis of the resulting problem is beyond the scope of our paper, our result may be seen as a useful first step in this direction.⁴⁶

Full support

Our results rely on a full support assumption for the distribution of preferences. Here we present an example showing that the characterization for ordinal preferences may not hold if one relaxes the assumption. Furthermore, we provide a natural refinement of OE that allows us to dispense with the assumption altogether.

EXAMPLE 4. Assume that there are two dates and two objects to be allocated at every date, $O_1 = O_2 = \{1, 2\}$. The set of pure allocations is $\mathbf{O} = \{(11), (12), (21), (22)\}$. Define the ordinal preference profiles

- $\pi_1(11) < \pi_1(12) < \pi_1(21) < \pi_1(22)$
- $\pi_2(12) < \pi_2(21) < \pi_2(11) < \pi_2(22)$
- $\pi_3(21) < \pi_3(12) < \pi_3(11) < \pi_3(22)$
- $\pi_4(22) < \pi_4(12) < \pi_4(21) < \pi_4(11)$.

Suppose that $F(\pi_k) > 0$ for $k = 1, 2, 3, 4$ and $F(\pi) = 0$ for $\pi \in \Pi \setminus \{\pi_1, \pi_2, \pi_3, \pi_4\}$, so that the full support assumption is violated. Define the cutoffs for each pure allocation as $\alpha_{11} = \alpha_{22} = 0$ and $\alpha_{12} = \alpha_{21} = 1$. Clearly, the cutoffs are not linear, as $\alpha_{11} + \alpha_{22} \neq \alpha_{12} + \alpha_{21}$. With budgets drawn uniformly in $[0, 1]$, a lottery-plus-cutoff mechanism \mathbf{x} would generate the random allocations $\mathbf{x}_{(11)}(\pi_1) = \mathbf{x}_{(11)}(\pi_2) = \mathbf{x}_{(11)}(\pi_3) = 1$ and $\mathbf{x}_{(22)}(\pi_4) = 1$. As the allocation rule \mathbf{x} is a lottery-plus-cutoffs mechanism, it is IC. It is easily checked that the random allocation is also OE if only agents π_1, π_2, π_3 , and π_4 are present in the market. Indeed, both π_1 and π_4 surely obtain their top choices, and agents π_2 and π_3 both obtain (11) so that improving mass transfers for either π_2 or π_3 would “hurt” π_4 . Hence, \mathbf{x} is an IC and OE mechanism, but it cannot be implemented by a GLC mechanism with linear cutoffs and, hence, with a spot mechanism.⁴⁷ Finally, note that if one reverses (11) and (22) in ranking π_2 —call this new preference ranking π_0 —then the resulting allocation would not be OE. Hence, if F assigns positive probability to π_0 (in addition to π_k for $k = 1, 2, 3, 4$), the issue is resolved.⁴⁸ \diamond

⁴⁶To capture the setting where capacities are hard constraints, one can simply let $C_{i,t}$ be infinite when $q_{i,t}$ is above a certain capacity constraint and 0 otherwise. Hence, our characterization result can be used even in a context where there are hard capacity constraints.

⁴⁷Indeed, to reproduce the random allocation \mathbf{x} , we would need that allocations (12) and (21) have strictly lower cutoffs than, respectively, (11) and (22). One can check that this cannot be achieved by linear prices since the ordering is similar to that in Example 1.

⁴⁸Example 4 demonstrates that our results are not merely the consequence of a duality argument implied by the assumption of a continuum of agents. This is an important difference with respect to the literature on linear pricing in combinatorial auctions, as we discussed in the [Introduction](#).

However, spot mechanisms can be seen as “robust” mechanisms once we relax the full support assumption. To see this, we refine OE and IC to get a full characterization without the full support assumption.

We say \mathbf{x} is *robustly OE and IC at F* if for all sequences⁴⁹ $F_n \rightarrow F$, there exists a sequence $\{\mathbf{x}_n\}$ such that $\mathbf{x}_n \rightarrow \mathbf{x}$, and \mathbf{x}_n is OE and IC at F_n for each n .^{50,51} To motivate the definition, consider the perspective of an analyst who does not know the precise distribution of preferences F . Then, to be confident in the mechanism \mathbf{x} , the analyst would want an allocation “close” to it to be OE and IC under the actual preference distribution. That is, the analyst would want \mathbf{x} to be robustly OE and IC. With this definition, we have the following proposition, the proof of which is relegated to Appendix D.

PROPOSITION 4. *Mechanism \mathbf{x} is robustly OE and IC at F if and only if it is a spot mechanism.*

Continuum

Our assumption of a continuum set of agents plays two roles in our environment. First, it is well known that ordinal efficiency and incentive compatibility (and equal treatment of equals) are incompatible in finite environments (see [Bogomolnaia and Moulin \(2001\)](#)). However, as shown in [Che and Kojima \(2010\)](#), there are incentive compatible mechanisms (respecting an equal treatment of equals) under which inefficiencies vanish when the market is large. These mechanisms are ordinally efficient in a continuum economy. Thus, the assumption of a continuum set of agents allows us to circumvent these impossibility results and to characterize ordinally efficient and incentive compatible allocation rules. Of course, these mechanisms may be inefficient in finite markets, but as argued in [Che and Kojima \(2010\)](#), inefficiencies vanish when the market grows large. Second, the continuum assumption allows us to avoid the standard issue in assignment problems with bundles that feasible fractional allocations may not correspond to lotteries over feasible integral allocations. Indeed, the Birkhoff–von Neumann theorem may fail in this environment (see [Budish \(2011\)](#) and [Nguyen, Peivandi, and Vohra \(2016\)](#) for a discussion of this issue). This issue does not arise in an economy with a continuum set of agents. The use of a continuum of agents to simplify the analysis of matching problems is not new (see, for instance, [Azevedo and Leshno \(2016\)](#); [Arnosti and Shi \(2020\)](#)). These models can usually be seen as the appropriate limit of a large finite problem.

Incentive compatibility

Theorems 3 and 4 both impose an efficiency notion (respectively, ordinal and cardinal) and IC. If one drops IC and only requires efficiency, then none of these theorems holds.

⁴⁹We endow the space of real numbers with a standard topology (e.g., induced by the Euclidean norm); the space of distributions is endowed with the topology of weak convergence and the product spaces are endowed with the product topology.

⁵⁰We say that \mathbf{x} is OE at F if our notion of OE holds under probability distribution F . We say that it satisfies IC at F if condition (4.1) for IC holds for any π that receives strictly positive mass under F .

⁵¹In this definition, one can replace “for all sequences $F_n \rightarrow F$ ” by “for some sequences of full-support distributions $F_n \rightarrow F$ ”; the next proposition still holds.

For instance, consider an economy with $T = 1$ and at least three objects. Let us consider an allocation rule that assigns to all preference rankings π but one their favorite pure allocation. Call π_0 the only preference ranking not obtaining its favorite pure allocation. Clearly, this allocation is ordinally efficient: whenever we try to improve the situation of π_0 , some other agents will be worse off. Further, the allocation rule is not incentive compatible, since π_0 , by swapping two objects in his or her ranking that differ from his or her top choice, will be guaranteed to get his or her top choice object. Hence, this allocation rule cannot be achieved by a generalized lottery-plus-cutoff mechanism (since these rules are incentive compatible).⁵²

Preference uncertainty

We assumed that the preferences of the agents over dynamic allocations are fixed at time zero. This implies that agents perfectly know their preferences when choosing an object at a given date under a spot mechanism. In particular, we do not allow agents to experience any unexpected shocks in their preferences over time. While it is reasonable for our main motivating example of course assignment, it can be less so for applications such as teacher assignment, where the time horizon is longer. However, at an informal level, these shocks in preferences should reinforce the value of spot mechanisms. Indeed, spot mechanisms are also attractive because agents do not commit to future assignments and can “re-optimize” over time in case of preference shocks. Hence, an environment where preference shocks occur over time may make spot mechanisms even better. We find this an interesting perspective that we leave for future research.

APPENDIX A: EXAMPLE 5

In Example 5 below, we exhibit a random assignment that is OE and cannot be replicated by a lottery-plus-cutoff mechanism with linear cutoffs. This, in particular, implies that we need to allow non-uniform distributions of budgets in the definition of spot mechanisms to achieve all OE and IC allocation rules.

EXAMPLE 5. Let $T = 2$ and $O_1 = O_2 = \{1, 2\}$, and consider the following spot mechanism where $p_1^1 = 0.6$, $p_2^1 = 0$, $p_1^2 = 0.4$ and $p_2^2 = 0$. The cutoffs are summarized as follows:

Allocation	Cutoff
(11)	1
(12)	0.6
(21)	0.4
(22)	0

⁵²While this example is fairly straightforward, it makes clear that our results are conceptually different from second welfare theorems. The latter do not impose any incentive compatibility constraints. In particular, this shows that Theorem 4 on cardinal mechanisms is conceptually different from the second welfare theorem of Miralles and Pycia (2020), who showed that in a finite market, efficient allocations with bundles can be implemented by competitive equilibria with linear prices.

Distribution G over possible budgets in $[0, 1]$ is assumed to satisfy $P(z = 1) = 0.2$, $P(z = 0.6) = 0.2$, $P(z = 0.4) = 0.1$, and $P(z = 0) = 0.5$. By Theorem 1, this random allocation is ordinally efficient. Now we claim that this random allocation cannot be replicated by a lottery-plus-cutoff mechanism with linear cutoffs. First, to replicate this allocation, it is clear that the order of cutoffs must remain the same, i.e., $\alpha_{11} > \alpha_{12} > \alpha_{21} > \alpha_{22}$. Given that, by definition of a lottery-plus-cutoff mechanism, the distribution over budgets in $[0, 1]$ must be uniform, we must have the following cutoffs to replicate the previous random allocation calculated with distribution G :

Allocation	Cutoff
(11)	0.8
(12)	0.6
(21)	0.5
(22)	0

However, it is easily checked that these cutoffs are nonlinear.⁵³ To recap, we need to use spot mechanisms with non-uniform distributions to reproduce the above OE random allocation rule. \diamond

APPENDIX B: PROVING THEOREM 3

The proof relies on the following result from the theory of linear inequalities.

LEMMA 3 (Carver, 1921). *For an arbitrary matrix \mathbf{A} , $\mathbf{Ax} < 0$ is feasible if and only if $\mathbf{y} = 0$ is the only solution for $\mathbf{y} \geq 0$ and $\mathbf{A}^T\mathbf{y} = 0$.*

To apply the lemma, we need some additional notation and preliminary results.

First, we discuss how the feasibility of linear system of inequalities relates to our notion of a GLC mechanism with linear cutoffs. We begin by describing bundles by vectors. Each feasible bundle $b \in B$ is assigned a row vector \mathbf{d}_b with $|O|$ columns, one column for each generalized object. For each object $i \in O$, we let $d_{bi} = 1$ if $i \in b$, and $d_{bi} = 0$ otherwise. It is useful to describe the differences in a composition between bundles b and b' by another row vector $\mathbf{a}_{b,b'}$ given by

$$\mathbf{a}_{b,b'} = \mathbf{d}_b - \mathbf{d}_{b'}.$$

Hence, each vector $\mathbf{a}_{b,b'}$ is composed only of 1s, -1 s, and 0s:

- If $i \in b$ and $i \notin b'$, then the row of $\mathbf{a}_{b,b'}$ corresponding to object i is equal to 1.
- If $i \notin b$ and $i \in b'$, then the row of $\mathbf{a}_{b,b'}$ corresponding to object i is equal to -1 .
- If i either belongs or does not belong to both bundles, then the row of $\mathbf{a}_{b,b'}$ corresponding to object i is equal to 0.

⁵³To see that these cutoffs are nonlinear, we need to argue that there is no vector $\mathbf{p} = (p_1^1, p_2^1, p_1^2, p_2^2)^T$ such that $\alpha_{ij} = p_i^1 + p_j^2$ for all $i, j = 1, 2$. Note that these equalities for $ij = 11, 12$ imply that $p_1^2 - p_2^2 = 0.2$, while equalities for $ij = 21, 22$ imply that $p_1^2 - p_2^2 = 0.5$.

For any order \leq on the set of feasible bundles B , we associate a matrix \mathbf{A} that captures the differences in a composition between each pair of strictly ordered bundles. In particular, let matrix \mathbf{A} contain as a row the vector $\mathbf{a}_{b,b'}$ as described above if and only if $b < b'$. Each column of \mathbf{A} corresponds to a generalized object. Let \mathbf{a}^i be the column of \mathbf{A} corresponding to object i . The following property of matrix \mathbf{A} is instrumental for the proof.

LEMMA 4. *For cutoffs α , let \mathbf{A} be the matrix associated with the total order on bundles induced by these cutoffs, i.e., $b < b' \Leftrightarrow \alpha_b < \alpha_{b'}$. Then there exist linear cutoffs $\bar{\alpha}$ such that for each $b, b' \in B$, $\alpha_b < \alpha_{b'} \Rightarrow \bar{\alpha}_b < \bar{\alpha}_{b'}$ if and only if there exists a vector \mathbf{p} such that $\mathbf{A}\mathbf{p} < 0$.*

PROOF. Let $\bar{\alpha}$ be linear cutoffs such that for each $b, b' \in B$, if $\alpha_b < \alpha_{b'}$, then $\bar{\alpha}_b < \bar{\alpha}_{b'}$, and let \mathbf{p} be the nonnegative prices inducing these cutoffs. Using our previous notations of vectors, it means that $\bar{\alpha}_b = \mathbf{d}_b \cdot \mathbf{p}$ for each $b \in B$.⁵⁴ The difference between cutoffs for any two bundles b and b' is $\bar{\alpha}_b - \bar{\alpha}_{b'} = \mathbf{a}_{b,b'} \cdot \mathbf{p}$. So, in particular, $\mathbf{a}_{b,b'} \cdot \mathbf{p} < 0$ means that bundle b has a lower cutoff than b' . Hence, $\mathbf{A}\mathbf{p} < 0$ because matrix \mathbf{A} contains a row $\mathbf{a}_{b,b'}$ if and only if $\alpha_b < \alpha_{b'}$.

Now suppose that there exists a vector \mathbf{p} such that $\mathbf{A}\mathbf{p} < 0$. Note that this vector can be arbitrary; in particular, having negative coordinates. We begin by showing that there exists a nonnegative price vector \mathbf{p}' such that $\mathbf{A}\mathbf{p}' < 0$ and the linear cutoffs induced by \mathbf{p}' belong to the unit interval as in the definition of the GLC mechanism, i.e., $\bar{\alpha}_b \leq 1$ for each $b \in B$. Note that each vector $\mathbf{a}_{b,b'}$ is composed of an equal number of 1s and -1 s because we have assumed that the bundles are of equal size. Hence $\mathbf{A}\mathbf{1} = 0$, where $\mathbf{1}$ is the unit vector. Therefore, for a sufficiently large $c > 0$, $\mathbf{p}'' = \mathbf{p} + c \cdot \mathbf{1}$ is a nonnegative price vector such that $\mathbf{A}\mathbf{p}'' < 0$. Moreover, for a sufficiently small $k > 0$, $\mathbf{p}' = k \cdot \mathbf{p}''$ is such that the linear cutoffs $\bar{\alpha}$ induced by \mathbf{p}' belong to the unit interval. Finally, by construction of \mathbf{A} , we have that for each $b, b' \in B$, if $\alpha_b < \alpha_{b'}$, then $\bar{\alpha}_b < \bar{\alpha}_{b'}$. \square

Second, we relate the existence of a nonnegative solution of a system of linear equations to the notion of ordinal efficiency. This requires restating the definition of ordinal efficiency in terms of probability mass transfers. Fix a random allocation \mathbf{q} , a preference ordering π , and a pair of bundles b and b' . Recall that q_b stands for the probability of getting bundle b . We say that $\tau_{b,b'}(\pi) \in \mathbb{R}$ is a *bilateral transfer from b to b' for π at \mathbf{q}* , or simply a bilateral transfer, if $0 < \tau_{b,b'}(\pi) \leq q_b$ and $q_{b'} + \tau_{b,b'}(\pi) \leq 1$. A bilateral transfer $\tau_{b,b'}(\pi)$ is *improving* if $\pi^{-1}(b') < \pi^{-1}(b)$. In words, an improving bilateral transfer $\tau_{b,b'}(\pi)$ specifies the probability mass is to be moved from a lower ranked bundle b to a higher ranked bundle b' . Now fix two random allocations \mathbf{q}' and \mathbf{q} . We say that \mathbf{q}' can be derived from \mathbf{q} by an *improving bilateral transfer* for π if there are bundles b and b' such that $q_{b''} = q'_{b''}$ for all bundles $b'' \in B \setminus \{b, b'\}$, and $q_b > 0$ and, moreover, $\tau_{b,b'}(\pi) := q_b - q'_b = q'_{b'} - q_{b'} > 0$ is an improving bilateral transfer from b to b' for π at \mathbf{q} . The following lemma applies the characterization of first-order stochastic dominance in terms of improving bilateral transfers to our framework.⁵⁵

⁵⁴We let the coordinates of \mathbf{p} and $\mathbf{a}_{b,b'}$ be ordered in such a way that vector operations make sense.

⁵⁵See, for instance, Østerdal (2010).

LEMMA 5. Fix a preference ordering π and two random allocations \mathbf{q} and \mathbf{q}' . The random allocation $\mathbf{q}' \neq \mathbf{q}$ first-order stochastically dominates \mathbf{q} for preferences π if and only if \mathbf{q}' can be derived from \mathbf{q} by a finite sequence of improving bilateral transfers. Formally, there exists a sequence $(\mathbf{q}_1, \dots, \mathbf{q}_n)$ of random allocations such that $\mathbf{q}_1 = \mathbf{q}$ and $\mathbf{q}_n = \mathbf{q}'$, and for $k = 1, \dots, n-1$, \mathbf{q}_{k+1} can be derived from \mathbf{q}_k by an improving bilateral transfer for π .

In light of Lemma 5, we can restate the second condition in the definition of ordinal efficiency. Specifically, for each preference π such that $\mathbf{x}(\pi) \neq \mathbf{x}'(\pi)$, we are required to find a sequence of improving bilateral transfers to go from random allocation $\mathbf{x}(\pi)$ to random allocation $\mathbf{x}'(\pi)$.

LEMMA 6. A random allocation \mathbf{x} is ordinally efficient if and only if there is no other random allocation \mathbf{x}' such that the following statements hold:

(i) For each object type $i \in O$, we have

$$\sum_{\pi \in \Pi} \sum_{b: i \in b} x'_b(\pi) F(\pi) = \sum_{\pi \in \Pi} \sum_{b: i \in b} x_b(\pi) F(\pi).$$

(ii) For each $\pi \in \Pi$ such that $\mathbf{x}(\pi) \neq \mathbf{x}'(\pi)$, random allocation $\mathbf{x}'(\pi)$ can be derived from $\mathbf{x}(\pi)$ by a sequence of improving bilateral transfers for π .

Now consider a vector \mathbf{y} whose coordinates are the same as those of the rows of matrix \mathbf{A} . We view each $y_{b,b'}$ as a probability mass to be transferred from a bundle b with a lower cutoff to a bundle b' with a higher cutoff.

LEMMA 7. Implementing the transfers in \mathbf{y} does not change the allocated mass of each object if and only if $\mathbf{A}^T \mathbf{y} = \mathbf{0}$.

PROOF. Consider a column vector \mathbf{y} , each coordinate of which, $y_{b,b'} \in \mathbb{R}$, corresponds to a row $\mathbf{a}_{b,b'}$ of \mathbf{A} . So \mathbf{y} specifies a set of probability mass transfers from lower to strictly higher bundles in the order of cutoffs. Now take a row i of matrix \mathbf{A}^T . Each coordinate of this row corresponds to some pair of bundles b and b' . For example, suppose b' has object i , while b does not. Then the corresponding coordinate of row i is equal to -1 . Imagine transferring mass $y_{b,b'}$ from b to b' . Then the total allocated mass of object i changes by $y_{b,b'}$. Therefore, the negative of the dot product of row i of \mathbf{A}^T and the vector \mathbf{y} , $-(\mathbf{a}^i)^T \mathbf{y}$, gives the total change in the allocated mass of object i resulting from the transfers defined by the vector \mathbf{y} . Accordingly, $-\mathbf{A}^T \mathbf{y}$ is a vector that captures the change in the allocated mass of each object. In particular, if $\mathbf{A}^T \mathbf{y} = \mathbf{0}$, then transfers \mathbf{y} simply redistribute the masses of objects across bundles. \square

Finally, we are ready to apply Lemma 3 to prove the following key technical result needed for the proof of Theorem 3.

PROPOSITION 5. Suppose that the distribution F has full support. Let \mathbf{x}^L be an ordinally efficient GLC mechanism with $L = (\boldsymbol{\alpha}, U_{[0,1]})$. Then there exist linear cutoffs $\bar{\alpha}$ such that for all $b, b' \in B$, if $\alpha_b < \alpha_{b'}$, then $\bar{\alpha}_b < \bar{\alpha}_{b'}$.

PROOF. Let \mathbf{A} be the matrix associated with the total order on bundles induced by cut-offs α . By Lemma 4, it suffices to show that there exists a vector \mathbf{p} such that $\mathbf{A}\mathbf{p} < \mathbf{0}$. For the sake of contradiction, suppose that such a vector does not exist. Then, by Lemma 3, there exists \mathbf{y} such that $\mathbf{y} \geq \mathbf{0}, \mathbf{y} \neq \mathbf{0}$ and $\mathbf{A}^T\mathbf{y} = \mathbf{0}$. Next we show that \mathbf{y} can be used to construct improving bilateral transfers for some preference profiles.

Let Π_α be the set of preference profiles whose ranking over bundles is consistent with the strict rankings induced by cutoffs α . Those are the preferences π such that for any $b, b' \in B$ with $\alpha_b < \alpha_{b'}$, then $\pi^{-1}(b') < \pi^{-1}(b)$. Below, we define a function f that, for each coordinate $y_{b,b'} > 0$ of \mathbf{y} , chooses a preference profile $f(b, b') \in \Pi_\alpha$ such that $x_b(f(b, b')) > 0$. This ensures that agents with preferences in $f(b, b')$ have a positive mass of b to transfer under an improving bilateral transfer from b to b' . Note that, by definition, profile $f(b, b')$ prefers bundle b' to b . For each $b \in B$, denote the set of bundles with a cutoff equal to α_b by $I(b) = \{b'' \neq b : \alpha_{b''} = \alpha_b\}$. Consider two cases:

- First, suppose $I(b) = \emptyset$. Then let $f(b, b')$ be any $\pi \in \Pi_\alpha$. Indeed, for all such π , we must have $x_b(\pi) > 0$ because a GLC mechanism with $L = (\alpha, U_{[0,1]})$ picks the budget of each agent uniformly from the unit interval. Hence, there is a positive probability for the event $E = \{\alpha_b \leq z < \hat{\alpha}_b\}$, where $\hat{\alpha}_b = \min\{\alpha_{b''} : \alpha_{b''} > \alpha_b\}$ is well defined since α_b is not the highest cutoff. Indeed, recall that \mathbf{y} contains coordinate $y_{b,b'}$ only when \mathbf{A} contains row $\mathbf{a}_{b,b'}$, which is true if and only if $\alpha_b < \alpha_{b'}$.
- Second, suppose $I(b) \neq \emptyset$. Then, by the full support assumption, there exists a preference profile $\pi_b \in \Pi_\alpha$ that ranks b ahead of each $b'' \in I(b)$. Hence, for the same reason as before, we must have $x_b(\pi_b) > 0$. So we define $f(b, b') = \pi_b$.

Now, for each $y_{b,b'} > 0$, pick the preference profile $\pi = f(b, b')$. By definition of a lottery-plus-cutoff mechanism, since $0 \leq \alpha_b < \alpha_{b'} \leq 1$, and because budgets are drawn uniformly in $[0, 1]$ and π prefers b' to b , we have that $x_{b'}(\pi) < 1$. For $\varepsilon > 0$, let all the agents with preferences π transfer a probability mass of $(\varepsilon/F(\pi))y_{b,b'}$ from b to b' at their random allocation $\mathbf{x}^L(\pi)$. Note that this is well defined given that, by the full support assumption, $F(\pi) > 0$ for all π . Hence, the total mass transferred from b to b' is $\varepsilon y_{b,b'} \geq 0$. Then, clearly, for a small enough $\varepsilon > 0$, these are improving bilateral transfers. Moreover, because $\mathbf{A}^T\mathbf{y} = \mathbf{0}$, by Lemma 4, these transfers do not change the allocated mass of each object. Therefore, \mathbf{x}^L is not ordinally efficient, which is a contradiction. It follows that there exist linear cutoffs $\bar{\alpha}$ such that for each $b, b' \in B$, if $\alpha_b < \alpha_{b'}$, then $\bar{\alpha}_b < \bar{\alpha}_{b'}$. \square

We are now in a position to prove Theorem 3.

PROOF OF THEOREM 3. (\Rightarrow) Let \mathbf{x} be an incentive compatible and ordinally efficient allocation rule. Our first step is to use Theorem 2 of Ashlagi and Shi (2016) to show that there exists a GLC mechanism that defines the same allocation rule as \mathbf{x} . This theorem applies to a static framework without bundles. However, we can reinterpret each bundle $b \in B$ as a single object and, hence, map our bundle framework back into a simple static environment. Specifically, interpret the set of feasible bundles B as a set of objects, so that preferences of agents over bundles can be thought of as preferences over

objects in a static setting. With this view, the model directly corresponds to the static case studied by Ashlagi and Shi (2016). Thus, \mathbf{x} is a well defined allocation rule in a static setting. An allocation rule \mathbf{x} is *AS-ordinally efficient* (AS-OE) if there is no other allocation rule \mathbf{x}' such that the following statements hold:

- (i) For each bundle $b \in B$, we have⁵⁶

$$\sum_{\pi \in \Pi} x'_b(\pi)F(\pi) = \sum_{\pi \in \Pi} x_b(\pi)F(\pi).$$

- (ii) For each $m = 1, \dots, |B|$ and each $\pi \in \Pi$, we have $\sum_{h=1}^m x'_{\pi(h)}(\pi) \geq \sum_{h=1}^m x_{\pi(h)}(\pi)$, with a strict inequality for some m and π such that $F(\pi) > 0$.

Using the same argument as in Lemma 1, it is immediate to see that any OE allocation rule must also be AS-OE. Hence, by Theorem 2 of Ashlagi and Shi (2016), we know that there exists a GLC mechanism with $\hat{L} = (\hat{\alpha}, U_{[0,1]})$ that defines the same allocation rule as \mathbf{x} , i.e., $\mathbf{x}^{\hat{L}} = \mathbf{x}$. Moreover, the corresponding GLC mechanism $L = (\alpha, U_{[0,1]})$ in our initial environment must also define the same allocation rule as \mathbf{x} , i.e., $\mathbf{x}^L = \mathbf{x}$.

Now, by Proposition 5, there exists a collection of linear cutoffs $\bar{\alpha}$ such that, for all $b, b' \in B$, if $\alpha_b < \alpha_{b'}$, then $\bar{\alpha}_b < \bar{\alpha}_{b'}$. However, using cutoffs $\bar{\alpha}$ while keeping the budget distribution $U_{[0,1]}$ will not generate the same allocation rule. Indeed, Proposition 5 only ensures that the linear cutoffs $\bar{\alpha}$ have the same strict ordering as α . In particular, absolute differences between cutoffs for any pair of bundles may change and bundles that had the same cutoffs in α may now have different cutoffs in $\bar{\alpha}$. To compensate for such changes, one has to adjust the budget distribution accordingly. In addition, we have to make sure that this adjustment can be done using a distribution with a continuous c.d.f. Let K be the number of distinct values taken by the cutoffs of α . First, index and relabel these unique values of cutoffs α to order them strictly so that $\alpha_1 < \alpha_2 < \dots < \alpha_K$. We partition bundles into equivalence classes: let $B_1 := \{b : 0 \leq \alpha_b \leq \alpha_1\}$ and, for $k = 2, \dots, K$, let $B_k := \{b : \alpha_{k-1} < \alpha_b \leq \alpha_k\}$. We also let $B_{\leq k} := \bigcup_{k'=1, \dots, k} B_{k'}$ be the set of all bundles affordable whenever the budget is equal to α_k . Note that under the GLC mechanism with $L = (\alpha, U_{[0,1]})$, the probability of being able to afford exactly all the bundles in $B_{\leq k}$ is $\alpha_{k+1} - \alpha_k$, where we set $\alpha_{K+1} := 1$. Fix a set B_k for some k . Under α , all the bundles in B_k have the same cutoff α_k . Under $\bar{\alpha}$, those bundles may not have the same cutoffs. We let $\bar{\alpha}_k^+ := \max_{b \in B_k} \bar{\alpha}_b$, $\bar{\alpha}_k^- := \min_{b \in B_k} \bar{\alpha}_b$, and $I_{k-1}^k := [\bar{\alpha}_{k-1}^+, \bar{\alpha}_k^-]$ for each $k \geq 1$, where we let $\bar{\alpha}_0^+ := 0$. By Lemma 5, because cutoffs $\bar{\alpha}$ have the same strict ordering as α , we have that $\bar{\alpha}_{k-1}^+ < \bar{\alpha}_k^-$ for $k \geq 1$ so that all the bundles in B_{k-1} have strictly lower cutoffs than those in B_k under $\bar{\alpha}$. Intuitively, going from cutoffs α to the linear cutoffs $\bar{\alpha}$ will create disjoint intervals of cutoffs $[\bar{\alpha}_k^-, \bar{\alpha}_k^+]$, one for each set B_k .⁵⁷ Since these intervals are disjoint, there are nonempty intervals between each of them, which are the intervals I_{k-1}^k defined above. To reproduce the allocation of the GLC mechanism with $L = (\alpha, U_{[0,1]})$

⁵⁶Remember that this is the key difference with our Condition (i) in the definition of ordinal efficiency. We impose equal mass for each object within a bundle, while Ashlagi and Shi (2016) impose equal mass for each bundle.

⁵⁷Note that these intervals of cutoffs can have a single point whenever $\bar{\alpha}_k^- = \bar{\alpha}_k^+$. It happens in particular when B_k is a singleton. This is illustrated in Figure 3 with the cutoff value α_3 .

using a GLC mechanism with $L' = (\bar{\alpha}, G)$ with a continuous c.d.f. G , we need to (i) put no probability mass in each interval $[\bar{\alpha}_k^-, \bar{\alpha}_k^+]$, and (ii) ensure that the probability of affording each bundle $B_{\leq k}$ is the same. Note that under $(\alpha, U_{[0,1]})$, this probability is $\alpha_{k+1} - \alpha_k$, while it is $G(\bar{\alpha}_{k+1}^-) - G(\bar{\alpha}_k^+)$ under the GLC mechanism with $L' = (\bar{\alpha}, G)$. To do so, let g be the probability density function (p.d.f.) defined as

$$g(x) = \begin{cases} \frac{\alpha_k}{\bar{\alpha}_k^- - \bar{\alpha}_{k-1}^+} & \text{if } x \in I_{k-1}^k \\ 0 & \text{otherwise.} \end{cases}$$

With this choice of p.d.f, it is immediate to see that the c.d.f. G is such that $G(\bar{\alpha}_k^-) = G(\bar{\alpha}_k^+) = \alpha_k$. Note that doing so ensures that g is a well defined probability density function. Intuitively, we put no probability mass in the intervals $[\bar{\alpha}_k^-, \bar{\alpha}_k^+]$ and we choose appropriately scaled uniform distributions for each interval I_{k-1}^k so that the resulting distribution is continuous and satisfies the requirements (i) and (ii) above. We illustrate our construction of the new c.d.f. G in Figure 3.

(\Leftarrow) Let \mathbf{x}^L be an allocation rule defined by a GLC mechanism with parameters $L = (\alpha, G)$ and linear cutoffs. We show that $\mathbf{x}^L \in \mathcal{M}_{IC}^e$. The incentive compatibility is straightforward, so we focus on proving ordinal efficiency. For the sake of contradiction, suppose \mathbf{x}^L is not ordinally efficient. Then there exists \mathbf{x}' such that \mathbf{x}^L and \mathbf{x}' allocate the same mass of each object and, for each π , the random allocation $\mathbf{x}'(\pi)$




FIGURE 3. Illustration of the transformed budget distribution. Note: Illustration of the construction of the new budget distribution. In the figure, we assumed that $\bar{\alpha}_3^- = \bar{\alpha}_3^+$ and we denoted $\bar{\alpha}_3$ their value. The 45 degree line is the c.d.f. of the uniform distribution for the GLC with $L = (\alpha, U_{[0,1]})$. The piecewise-linear function is the c.d.f. G of the equivalent GLC mechanism with $L' = (\bar{\alpha}, G)$ that uses linear cutoffs.

can be derived from $\mathbf{x}^L(\pi)$ via a sequence of improving bilateral transfers (whenever $\mathbf{x}'(\pi) \neq \mathbf{x}^L(\pi)$). Given such a sequence for π , let $\tau_{b,b'}(\pi)$ be the total mass transferred from bundle b to bundle b' .⁵⁸ We first note that if $\tau_{b,b'}(\pi) > 0$, then we must have that $\alpha_b < \alpha_{b'}$. Indeed, assume that $\alpha_{b'} \leq \alpha_b$. By definition of improving transfers, we must have that $\pi^{-1}(b') < \pi^{-1}(b)$ and whenever an agent with preferences π has budget $z \geq \alpha_b$, both b and b' can be chosen by this agent so that she always picks b' . Hence, it implies $x_b^L(\pi) = 0$, a contradiction to $\tau_{b,b'}(\pi)$ being the sum of the improving bilateral transfers from b to b' . Now we aggregate the bilateral transfers across all agents into a column vector \mathbf{y} . In particular, for each $b, b' \in B$ such that $\alpha_b < \alpha_{b'}$, we let

$$y_{b,b'} = \sum_{\pi \in \Pi} \tau_{b,b'}(\pi) F(\pi).$$

Hence, $y_{b,b'}$ is the total mass transferred by all agents from b to b' . Let \mathbf{A} be the matrix associated with the total order on bundles induced by cutoffs α . Because \mathbf{x}^L and \mathbf{x}' allocate the same mass of each object, by Lemma 4, we have $\mathbf{A}^T \mathbf{y} = 0$. In addition, since $\mathbf{x}^L \neq \mathbf{x}'$, by construction, we have $\mathbf{y} \neq 0$. But then, by Lemma 3, $\mathbf{A}\mathbf{p} < 0$ is not feasible, a contradiction to α being linear. Therefore, allocation rule \mathbf{x}^L is ordinally efficient. \square

APPENDIX C: PROVING THEOREM 4

The proof resembles the proof of Theorem 3. In particular, we will use a matrix \mathbf{A} associated with an order over bundles induced by the prices α of a CEEI to prove the existence of certain linear prices. However, the main difference is that we will not be using all the strict orderings induced by α .

Let \mathbf{x} be an incentive compatible and Pareto efficient allocation rule. First, we apply the same connection to Ashlagi and Shi (2016) as in the proof of Theorem 3. Hence, in this new environment, we treat bundles as objects. An allocation rule \mathbf{x} is AS-Pareto efficient if there is no other allocation rule \mathbf{x}' such that the following statements hold:

- (i) For each bundle $b \in B$, we have $\int_U x'_b(\mathbf{u}) dF = \int_U x_b(\mathbf{u}) dF$.
- (ii) For each $\mathbf{u} \in U$, we have $\mathbf{u} \cdot \mathbf{x}'(\mathbf{u}) \geq \mathbf{u} \cdot \mathbf{x}(\mathbf{u})$ and there is a set $A \subset U$ such that $F(A) > 0$ and the inequality is strict for each $\mathbf{u} \in A$.

Similarly to Lemma 1, it is immediate to see that any Pareto efficient allocation rule must also be AS-Pareto efficient.⁵⁹ Hence, by Theorem 1 of Ashlagi and Shi (2016), we know that, with a continuous distribution F with full relative support, the mechanism \mathbf{x} is a CEEI for some prices $(\alpha_b)_{b \in B} \in (0, \infty]^{|B|}$. Note that these prices, following the definition of Ashlagi and Shi (2016), are strictly positive and some of them can be infinite. We let

⁵⁸By definition of ordinal efficiency, one can ignore the bilateral transfers for agents π with $F(\pi) = 0$. This implies that this part of the proof (i.e., (\Leftarrow)) does not use the full ordinal support assumption.

⁵⁹Indeed, the allocation \mathbf{x}' in the definition of AS-Pareto efficiency would also be valid if one uses the definition of Pareto efficiency in Section 5.

$\alpha_{\max} := \max_{b \in B} \alpha_b$ and $\alpha_{\min} := \min_{b \in B} \alpha_b$. We start with the following simple observation.

LEMMA 8. *If \mathbf{x} is CEEI for prices α , then $\alpha_{\min} \leq 1$.*

PROOF. If $\alpha_b > 1$ for each bundle b , then for each $\mathbf{q} \in \Delta$, we have $\alpha \cdot \mathbf{q} > 1$. Hence, $\mathbf{x}(\mathbf{u})$ is not affordable for each \mathbf{u} , a contradiction. \square

Now fix new prices $\hat{\alpha} \in [0, \infty]^{|B|}$. Note that we now allow these prices to be null. We call an allocation rule \mathbf{x} a r -CEEI with prices $\hat{\alpha}$ if, intuitively, it is a CEEI with budget r (instead of 1), i.e., $\arg \max_{\mathbf{q} \in \Delta} \{\mathbf{u} \cdot \mathbf{q} : \hat{\alpha} \cdot \mathbf{q} \leq r\}$. As before, we can similarly define $\hat{\alpha}_{\min}$.

LEMMA 9. *If \mathbf{x} is a CEEI with prices $\alpha \in (0, \infty]^{|B|}$, then \mathbf{x} is an r -CEEI with prices $\hat{\alpha} \in [0, \infty]^{|B|}$, budget $r = 1 - \alpha_{\min}$, and $\hat{\alpha}_{\min} = 0$.*

PROOF. By Lemma 8, we have $\alpha_{\min} \leq 1$. Then reducing the budget and all the prices by α_{\min} does not change the budget set and, hence, \mathbf{x} is still a CEEI under the reduced budget and prices, i.e., an r -CEEI and prices $\hat{\alpha} \in [0, \infty]^{|B|}$ with $\hat{\alpha}_{\min} = 0$ as required. \square

Fix prices $\alpha \in [0, \infty]^{|B|}$. Let $B_{\infty} = \{b : \alpha_b = \infty\}$. Note that B_{∞} is nonempty when $\alpha_{\max} = \infty$. We denote the highest finite price by $\alpha_{\max^*} := \max_b \{\alpha_b : \alpha_b < \infty\}$, and let b^{\max^*} be a bundle such that $\alpha_b = \alpha_{\max^*}$. The following lemma is useful to construct improving bilateral transfers.

LEMMA 10. *If \mathbf{x} is an r -CEEI with prices $\alpha \in [0, \infty]^{|B|}$ and $\alpha_{\min} = 0$, let the set $P \subset B^2$ be defined as follows:*

- *Case 1. If $\alpha_{\max^*} \leq r$, let $P := B \setminus B_{\infty} \times B_{\infty}$ with $B_{\infty} := \{b : \alpha_b = \infty\}$.*
- *Case 2. If $\alpha_{\max^*} > r$ and $r = 0$, let $B_0 := \{b : \alpha_b = 0\}$ and $P := B_0 \times B \setminus B_0$.*
- *Case 3. If $\alpha_{\max^*} > r$ and $r > 0$, let $P := \{(b, b') : \alpha_b < \alpha_{b'}\}$.*

Then, for any pair $(b, b') \in P$, the following statements hold:

- (a) *We have $\alpha_b < \alpha_{b'}$.*
- (b) *There is an open set $f(b, b') \subset U$ such that (i) $u_b < u_{b'}$ and (ii) for some $m > 0$, $x_b(\mathbf{u}) \geq m$ for all $\mathbf{u} \in f(b, b')$.*

PROOF. As discussed above, if \mathbf{x} is an incentive compatible and Pareto efficient allocation rule, then, using Lemma 9, it is an r -CEEI with prices $\alpha \in [0, \infty]^{|B|}$ such that $\alpha_{\min} = 0$. We follow each case of the lemma.

Case 1: $\alpha_{\max^} \leq r$.* In this case, all bundles with a finite price are affordable. By definition, for each $(b, b') \in P$, we have $\alpha_b < \alpha_{b'}$ so the first condition of the lemma holds.

For each $(b, b') \in P$, let $f(b, b') \subset U$ be the set of utility vectors such that

- $u_{b'} = 2M + \varepsilon_{b'}$ with $\varepsilon_{b'} \in (0, \bar{\varepsilon})$

- $u_b = M + \varepsilon_b$ with $\varepsilon_b \in (0, \bar{\varepsilon})$
- $u_{b''} = \varepsilon_{b''}$ with $\varepsilon_{b''} \in (0, \bar{\varepsilon})$ for each $b'' \neq b, b'$,

where M and $\bar{\varepsilon}$ are some constants. Clearly, the set $f(b, b')$ is open in U as a product of open intervals. For $M > \bar{\varepsilon}$, bundle b' gives the highest utility followed by b , followed by all other bundles. Note that since $(b, b') \in B \setminus B_\infty \times B_\infty$ and $\alpha_{\max^*} \leq r$, b is always affordable under the CEEL, while b' is not. Since b gives the highest utility among affordable bundles for $\mathbf{u} \in f(b, b')$, we have $x_b(\mathbf{u}) = 1$ as required.

Case 2: $\alpha_{\max^} > r$ and $r = 0$.* In this case, only free bundles are affordable. Since $\alpha_{\min} = 0$, the set B_0 is nonempty. By construction, we have $\alpha_b = 0 < \alpha_{b'}$ for any $(b, b') \in P$ so that the first requirement of the lemma holds.

For each $(b, b') \in P$, we now define the set $f(b, b') \subset U$ in the same way as in Case 1. Note that since $(b, b') \in B_0 \times B \setminus B_0$ and $r = 0$, bundle b is always affordable under the CEEL, while b' is not. Since b gives the highest utility among affordable bundles for $\mathbf{u} \in f(b, b')$, we have $x_b(\mathbf{u}) = 1$ as required.

Case 3: $\alpha_{\max^} > r$ and $r > 0$.* In this case, note that we can normalize the budget to one by dividing all the prices by r and obtain the same CEEL. So, in what follows, we assume that $r = 1$. Remember that we have at least one free bundle since $\alpha_{\min} = 0$. By definition of P , the first requirement of the lemma holds.

Fix constants M , $\bar{\varepsilon}$, and $\bar{\delta}$. Remember that b^{\max^*} is a bundle such that $\alpha_{b^{\max^*}} = \alpha_{\max^*}$. For each $(b, b') \in P$, let $f(b, b')$ be the set of utility vectors such that

- $u_{b''} = \alpha_{b''} + \varepsilon_{b''}$ with $\varepsilon_{b''} \in (0, \bar{\varepsilon})$ for $b'' \neq b, b^{\max^*}$ such that $\alpha_{b''} < \infty$
- $u_{b''} = M + \varepsilon_{b''}$ with $\varepsilon \in (0, \bar{\varepsilon})$ for $b'' \in B_\infty$
- $u_b = \alpha_b + \delta_b + \varepsilon_b$ with $\varepsilon_b \in (0, \bar{\varepsilon})$
- $u_{b^{\max^*}} = \alpha_{b^{\max^*}} + \delta_{b^{\max^*}} + \varepsilon_{b^{\max^*}}$ with $\varepsilon_{b^{\max^*}} \in (0, \bar{\varepsilon})$.

In words, utility vectors in $f(b, b')$ assign to each bundle b'' a utility equal to the bundle's price $\alpha_{b''}$ (or a large constant if this price is infinite) perturbed by some positive constant. For each bundle $b'' \neq b^{\max^*}$ with price $\alpha_{b''} < \infty$, let $s(\alpha_{b''})$ be the next strictly highest price, possibly infinite, i.e., $s(\alpha_{b''}) := \min_{b'} \{\alpha_{b'} : \alpha_{b'} > \alpha_{b''}\}$. We can choose positive constants M , δ_b , $\delta_{b^{\max^*}}$, and $\bar{\varepsilon}$, so that they satisfy the following constraints:

- (i) For each b'' such that $\alpha_{b''} \neq \alpha_{b^{\max^*}}$, we have

$$\alpha_{b''} + \delta_b + \bar{\varepsilon} < s(\alpha_{b''}) \quad (\text{C.1})$$

and $M > \alpha_{b^{\max^*}} + \delta_{b^{\max^*}} + \bar{\varepsilon}$.

- (ii) We have

$$\delta_b > \bar{\varepsilon}. \quad (\text{C.2})$$

(iii) If $\alpha_b > 0$ and $b \neq b^{\max^*}$, then for each $b'' \neq b, b^{\max^*}$ such that $\alpha_{b''} > 0$, we have

$$\frac{\delta_b}{\alpha_b} > \frac{\delta_{b^{\max^*}} + \bar{\varepsilon}}{\alpha_{b^{\max^*}}} + \left(\frac{1}{\alpha_b} - \frac{1}{\alpha_{b^{\max^*}}} \right) \bar{\varepsilon} > \frac{\delta_{b^{\max^*}}}{\alpha_{b^{\max^*}}} > \frac{\bar{\varepsilon}}{\alpha_{b''}}. \tag{C.3}$$

The constraint (C.1) makes sure that the ranking induced by the perturbed utilities is consistent with the strict ranking induced by prices α . Constraint (C.2) implies that bundle b is the most attractive bundle among all bundles with the same price. Constraint (C.3) implies that bundles b and b^{\max^*} deliver the highest utility per unit of artificial currency among all non-free bundles with finite price, and, roughly speaking, b is sufficiently more attractive than b^{\max^*} .⁶⁰ Clearly, the set $f(b, b')$ is open in U as a product of open intervals in \mathbb{R} .

We now show that if $\mathbf{u} \in f(b, b')$, then $x_b(\mathbf{u}) = m$ for some $m > 0$. We begin by showing that, in the CEEL, there does not exist $b'' \neq b, b^{\max^*}$ and $\mathbf{u} \in f(b, b')$ such that $\alpha_{b''} > 0$ and $x_{b''}(\mathbf{u}) > 0$. For the sake of contradiction, suppose such b'' and \mathbf{u} exist. Consider reducing expenditures of such agents on b'' by $\eta > 0$ and increasing their expenditures on b^{\max^*} by η so their probability share of b'' decreases by $\eta/\alpha_{b''}$ and their probability share of b^{\max^*} increases by $\eta/\alpha_{b^{\max^*}} \leq \eta/\alpha_{b''}$. To keep the sum of probability shares equal to 1, increase the share of any free bundle by $\eta/\alpha_{b''} - (\eta/b^{\max^*})$. For a sufficiently small $\eta > 0$, such transfer of mass is feasible and increases the utility of agents with $\mathbf{u} \in f(b, b')$ by constraint (C.1) above, a contradiction to the allocation being a CEEL.

First, suppose $\alpha_b = 0$. Then, given the above result, an agent with $\mathbf{u} \in f(b, b')$ must spend her entire budget on b^{\max^*} in purchasing a $1/\alpha_{b^{\max^*}} < 1$ probability share of b^{\max^*} , and complete the allocation with the free bundle b in purchasing a $1 - (1/\alpha_{b^{\max^*}}) > 0$ probability share of b , because $\delta_b > \bar{\varepsilon}_{b''}$ for each $b'' \neq b$ such that $\alpha_{b''} = 0$.

Second, suppose $\alpha_b > 0$. Notice that because of constraint (C.3), an agent with $\mathbf{u} \in f(b, b')$ must allocate the entire budget between bundles b and b^{\max^*} , and potentially complete the allocation with a share of a free bundle that delivers the highest utility, denoted by b_0 . Specifically, she solves the optimization problem

$$\max_{0 \leq z \leq 1} (\alpha_{b^{\max^*}} + \delta_{b^{\max^*}} + \varepsilon_{b^{\max^*}}) \frac{z}{\alpha_{b^{\max^*}}} + (\alpha_b + \delta_b + \varepsilon_b) \frac{1-z}{\alpha_b} + \varepsilon_{b_0} \left(1 - \frac{z}{\alpha_{b^{\max^*}}} - \frac{1-z}{\alpha_b} \right)$$

subject to

$$1 - \frac{z}{\alpha_{b^{\max^*}}} - \frac{1-z}{\alpha_b} \geq 0.$$

Given the constraint (C.3), the objective is linearly decreasing in z . If $\alpha_b \geq 1$, then the constraint does not bind and optimally $x_b(\mathbf{u}) = 1/\alpha_b$, i.e., the entire budget is spent on

⁶⁰Note that there are positive constants $M, \delta_b, \delta_{b^{\max^*}}$, and $\bar{\varepsilon}$ satisfying (i), (ii), and (iii). Indeed, one can set δ_b and $\bar{\varepsilon}$ small enough and M high enough so that (i) holds. With an even smaller $\bar{\varepsilon}$, (ii) holds. Finally, with $\bar{\varepsilon}$ small again and $\delta_{b^{\max^*}}$ small, (iii) holds true. Also note that since we have assumed that $\alpha_{\max^*} < \infty$, then $\alpha_{b^{\max^*}} < \infty$, so that constraint (C.3) is indeed true.

b . If $\alpha_b < 1$, then the constraint binds, which implies that $x_{b_0}(\mathbf{u}) = 0$, and the budget is split between b^{\max^*} and b such that

$$x_b(\mathbf{u}) = \frac{\alpha_{b^{\max^*}}(1 - \alpha_b)}{\alpha_{b^{\max^*}} - \alpha_b}.$$

Summarizing, for each $\mathbf{u} \in f(b, b')$, we have

$$x_b(\mathbf{u}) \geq \min \left\{ 1 - \frac{1}{\alpha_{b^{\max^*}}}, \frac{1}{\alpha_b}, \frac{\alpha_{b^{\max^*}}(1 - \alpha_b)}{\alpha_{b^{\max^*}} - \alpha_b} \right\} := m > 0$$

as required. \square

Last, we show that from the sets $f(b, b')$ as in Lemma 10, we can create an open cone of preferences with positive mass having the same property. The reader interested in the proof of Theorem 4 can skip the proof since it mostly relies on topology arguments.

LEMMA 11. *Fix an r -CEEI \mathbf{x} with prices α . If the distribution F has full relative support, then for each pair $(b, b') \in P$ and its associated open set $f(b, b')$ from Lemma 10, there exists an open cone $C(b, b') \in \mathcal{C}$ such that $F(C(b, b')) > 0$ and for each $\mathbf{u} \in C(b, b')$ and some $m > 0$, we have $x_b(\mathbf{u}) \geq m$.*

PROOF. Fix a pair $(b, b') \in P$ and the associated open set $f(b, b') \subset U$ from Lemma 10.⁶¹ In the sequel, we recall that Proj_D stands for the projection from U into D , i.e.,

$$\text{Proj}_D(\mathbf{u}) := \left(u_b - \frac{\sum u_b}{|B|} \right)_b.$$

Note that, from Lemma 10, under an r -CEEI, for any $\mathbf{u} \in f(b, b')$, then $x_b(\mathbf{u}') \geq m > 0$. For any $\mathbf{u}' = \lambda \mathbf{u} - \xi \mathbf{1}$ with $\lambda > 0$ and $\xi \in \mathbb{R}$, since the choices are invariant to linear transformations of \mathbf{u} , we also have $x_b(\mathbf{u}') \geq m$. In words, rescaling and translating the cardinal utilities will not impact the optimal choice of the agent in a CEEI. Given $\lambda > 0$, we denote $X_\lambda := \{\mathbf{u}' \in U : \mathbf{u}' = \lambda \mathbf{u} \text{ for some } \mathbf{u} \in f(b, b')\}$. Note that for any $\lambda > 0$, X_λ is open in U (since the function $\mathbf{u} \mapsto \lambda \mathbf{u}$ is a homeomorphism). Now let us consider $\mathcal{Z} := \bigcup_{\lambda > 0} X_\lambda$. Note that, as a union of open sets, \mathcal{Z} is open in U . Let $C := \text{Proj}_D(\mathcal{Z})$. Here again, for any $\mathbf{u} \in C$, we must have $x_b(\mathbf{u}) \geq m$, since such \mathbf{u} are simple linear transformations of utility vectors in $f(b, b')$.

We first claim that C is a cone. Take any $\mathbf{u}' \in C$ and any $\lambda > 0$. We must show that $\lambda \mathbf{u}' \in C$. Indeed, since $\mathbf{u}' \in C$, we must have that for some $\mathbf{u} \in \mathcal{Z}$, $\text{Proj}_D(\mathbf{u}) = \mathbf{u}'$. Hence, $\text{Proj}_D(\lambda \mathbf{u}) = \lambda \text{Proj}_D(\mathbf{u}) = \lambda \mathbf{u}'$, where the first equality uses the linearity of Proj_D . Since, by definition of set \mathcal{Z} , it must be that $\lambda \mathbf{u}$ belongs to \mathcal{Z} , $\text{Proj}_D(\lambda \mathbf{u}) = \lambda \mathbf{u}'$ implies that $\lambda \mathbf{u}' \in \text{Proj}_D(\mathcal{Z}) = C$, as claimed.

Now, we show that C is open in D so as to eventually show that C is open in \mathcal{C} . This comes from the feature that Proj_D is an open map together with the fact that \mathcal{Z} is open in

⁶¹Remember that the sets P and $f(b, b')$ change, depending on the values of the cutoffs α as shown in the proof of Lemma 10.

U .⁶² Finally, we want to show that our cone C is open in \mathcal{C} , i.e., $C \cap \tilde{D}$ is open in \tilde{D} . This is true since, as we just claimed, C is open in D and so $C \cap \tilde{D}$ is open in \tilde{D} by definition of the relative topology. Thus, we can set $C(b, b') := C$. The open cone $C(b, b')$ satisfies $x_b(\mathbf{u}) \geq m$ for any $\mathbf{u} \in C(b, b')$. Since F has full relative support and $C(b, b')$ is open in \tilde{D} , we have $F(C(b, b')) > 0$. \square

We are now equipped with all the lemmas to prove Theorem 4. Similar to the proof of Theorem 3, the proof relies on constructing feasible bilateral transfers whenever there is no solution to a well constructed set of linear inequalities. The proof is divided into several cases, depending on the value of α_{\max} and α_{\min} . For each of them, we will show that we can construct a spot MRB that induces the same allocation as the initial CEEI. As in the proof of Proposition 5, we will construct a matrix \mathbf{A} associated to a strict ordering over bundles in B and consider the system of linear inequalities $\mathbf{A}\mathbf{p} < 0$. We will show that if such a system has no solution, there exists \mathbf{y} such that $\mathbf{y} \geq \mathbf{0}$, $\mathbf{y} \neq \mathbf{0}$, and $\mathbf{A}^T \mathbf{y} = 0$, and that such \mathbf{y} can be used to construct feasible improving bilateral transfers for a positive mass of agents (using Lemma 11) so that \mathbf{x} is not Pareto efficient.

PROOF OF THEOREM 4. Suppose \mathbf{x} is an incentive compatible and Pareto efficient allocation rule. By Lemma 9, \mathbf{x} is an r -CEEI with prices $\alpha \in [0, \infty]^{|B|}$ such that $\alpha_{\min} = 0$. Fix the set P of pairs of bundles as defined in Lemma 10. First, we show that there exist linear prices $\bar{\alpha}$ such that $\bar{\alpha}_b < \bar{\alpha}_{b'}$ for each $(b, b') \in P$.

Construct the matrix \mathbf{A} as in the proof of Proposition 5 so that each row of \mathbf{A} corresponds to a pair of bundles $(b, b') \in P$ and each column corresponds to a generalized object in O . For the sake of contradiction suppose such linear prices $\bar{\alpha}$ do not exist. Then, as in Proposition 5, there exists \mathbf{y} such that $\mathbf{y} \geq \mathbf{0}$, $\mathbf{y} \neq \mathbf{0}$, and $\mathbf{A}^T \mathbf{y} = 0$, and, in what follows, we use \mathbf{y} to construct improving bilateral transfers for a positive mass of agents.

For each $(b, b') \in P$, Lemma 11 guarantees that there exists an open cone $C(b, b')$ such that $F(C(b, b')) > 0$, and for each $\mathbf{u} \in C(b, b')$, we have $u_b < u_{b'}$ and $x_b(\mathbf{u}) \geq m > 0$. For each $(b, b') \in P$, consider a transfer of a probability mass $(\epsilon/F(C(b, b'))y_{b,b'})$ from b to b' for agents with $\mathbf{u} \in C(b, b')$ at their random allocation $\mathbf{x}(\mathbf{u})$. By construction of $C(b, b')$, this is an improving bilateral transfer given a sufficiently small $\epsilon > 0$. Moreover, because $\mathbf{A}^T \mathbf{y} = 0$, by Lemma 7, these transfers do not change the allocated mass of each object. Therefore, \mathbf{x} is not Pareto efficient, which is a contradiction. It follows that there exist linear cutoffs $\bar{\alpha}$ such that $\bar{\alpha}_b < \bar{\alpha}_{b'}$ for each $(b, b') \in P$. Without loss of generality assume that $\max_b \bar{\alpha}_b = 1$.

To finish the proof, we construct a spot MRB mechanism that implements r -CEEI \mathbf{x} . By Proposition 3, \mathbf{x} is a MRB mechanism $L = (\hat{\alpha}, \mathcal{G})$, where cutoffs $\hat{\alpha}$ are the normalized prices α as in the proof of Proposition 3. Using the linear prices $\bar{\alpha}$, we now construct a collection of distributions \mathcal{G}' such that the spot MRB mechanism with $L' = (\bar{\alpha}, \mathcal{G}')$ implements the allocation rule \mathbf{x} . For each distribution $G_{\mathbf{x}(\mathbf{u})} \in \mathcal{G}$, let the corresponding

⁶²Projection Proj_D is a continuous mapping under our topologies, and it is surjective and linear. By the open mapping theorem, Proj_D is an open mapping, i.e., for any open set \mathcal{O} in U , $\text{Proj}_D(\mathcal{O})$ is open in D .

distribution $G'_{\mathbf{x}(\mathbf{u})} \in \mathcal{G}'$ assign probability $x_b(\mathbf{u})$ to $\bar{\alpha}_b$ instead of $\hat{\alpha}_b$. We now consider the same cases as in Lemma 10.

Case 1: $\alpha_{\max^*} \leq r$. In the r -CEEI, each agent is assigned her favorite bundle out of those with a finite price. In the spot MRB with $L' = (\bar{\alpha}, \mathcal{G}')$, each agent can also receive the same bundle and cannot receive a positive share of any bundle with an infinite price because those bundles keep having the highest prices under the linear cutoffs $\bar{\alpha}$ and the distributions in \mathcal{G}' put probability 1 on the budget strictly below these prices. Hence, the induced allocation rule must be the same.

Case 2: $\alpha_{\max^*} > r$ and $r = 0$. In the r -CEEI, each agent is assigned her favorite free bundle. Similar to the previous case, each agent can also receive the same bundle in the spot MRB with $L' = (\bar{\alpha}, \mathcal{G}')$. Moreover, she cannot receive a positive share of any other bundle because those bundles have strictly higher prices than the prices of the free bundles under the linear cutoffs $\bar{\alpha}$ and the distributions in \mathcal{G}' put probability 1 on the budget strictly below these prices. Hence, the induced allocation rule must be the same.

Case 3: $\alpha_{\max^*} > r$ and $r > 0$. Note that in $L' = (\bar{\alpha}, \mathcal{G}')$, for each realization of a random budget, the set of affordable bundles is the same as in L for each distribution because the linear cutoffs $\bar{\alpha}$ have the same strict order as prices α . Then we have that, for each distribution and for each set of bundles, the probability that this set is affordable is the same in L and L' . Hence, the induced allocation rule must be the same. \square

APPENDIX D: PROVING PROPOSITION 4

Recall that a spot mechanism \mathbf{x} is characterized by a GLC with parameters (α, G) , where $x_{\pi(h)}(\pi) = G(\min_{m=1, \dots, h-1} \alpha_{\pi(m)}) - G(\min_{m=1, \dots, h} \alpha_{\pi(m)})$ for every π and $h = 1, \dots, |\mathbf{O}|$. In addition, we know that there exists nonlinear $\mathbf{p} = (p^t)_{t=1, \dots, T}$, where $p^t = (p^t_i)_{i \in O_t}$ for each $t = 1, \dots, T$ satisfying

$$\alpha_{\mathbf{o}} = \sum_{t=1}^T p^t_{o_t}$$

for each $\mathbf{o} = (o_1, \dots, o_T) \in \mathbf{O}$. We say that (α, G, \mathbf{p}) corresponds to spot mechanism \mathbf{x} .

LEMMA 12. *Take a sequence $\mathbf{x}^n \rightarrow \mathbf{x}$, where, for each n , \mathbf{x}^n is a spot mechanism. Further, assume that the corresponding sequence $(\alpha^n, G^n, \mathbf{p}^n)$ converges to (α, G, \mathbf{p}) . We must have that \mathbf{x} is a spot mechanism and (α, G, \mathbf{p}) corresponds to \mathbf{x} .*

PROOF. Since for each n ,

$$x_{\pi(h)}^n(\pi) = G^n\left(\min_{m=1, \dots, h-1} \alpha_{\pi(m)}^n\right) - G^n\left(\min_{m=1, \dots, h} \alpha_{\pi(m)}^n\right)$$

for every π and $h = 1, \dots, |\mathbf{O}|$, the same must hold as well in the limit, i.e.,

$$x_{\pi(h)}(\pi) = G\left(\min_{m=1, \dots, h-1} \alpha_{\pi(m)}\right) - G\left(\min_{m=1, \dots, h} \alpha_{\pi(m)}\right)$$

for every π and $h = 1, \dots, |\mathbf{O}|$.

In addition, since for each n , for each $t = 1, \dots, T$, we have

$$\alpha_{\mathbf{o}}^n = \sum_{t=1}^T p_{o_t}^{t,n}$$

for each $\mathbf{o} = (o_1, \dots, o_T) \in \mathbf{O}$, the same must hold as well in the limit, i.e.,

$$\alpha_{\mathbf{o}} = \sum_{t=1}^T p_{o_t}^t.$$

Hence, \mathbf{x} is a spot mechanism and (α, G, \mathbf{p}) corresponds to \mathbf{x} , as claimed. \square

PROOF OF PROPOSITION 4. (\Rightarrow) Assume that \mathbf{x} is robustly OE and IC at F . Pick a sequence $F_n \rightarrow F$, where F_n has full support. Because \mathbf{x} is robustly OE and IC at F , we know that there is a sequence $\{\mathbf{x}_n\}$ such that $\mathbf{x}_n \rightarrow \mathbf{x}$, and \mathbf{x}_n is OE and IC at F_n for each n . Note that, by Theorem 1, this implies that \mathbf{x}_n is a spot mechanism for each n . Let $(\alpha^n, G^n, \mathbf{p}^n)$ correspond to \mathbf{x}_n for each n . Note that α^n and \mathbf{p}^n clearly lie in a (sequentially) compact set. In addition, the space of probability measures over the compact set $[0, 1]$ is sequentially compact in the topology of weak convergence of measures. So G^n also lies in a sequentially compact set.⁶³ Thus, taking a subsequence if necessary, we can assume that $(\alpha^n, G^n, \mathbf{p}^n) \rightarrow (\alpha, G, \mathbf{p})$. By Lemma 12, \mathbf{x} is a spot mechanism.

(\Leftarrow) Assume that \mathbf{x} is a spot mechanism. By Theorem 1 (and the observation that Theorem 1(\Leftarrow) holds without the full-support assumption; see footnote 58), \mathbf{x} is OE and IC for all distributions F' . Now fix any sequence of distributions $F_n \rightarrow F$ and let \mathbf{x}_n be the constant sequence equal to \mathbf{x} for all n . By the previous observation, \mathbf{x}_n is OE and IC at F_n for each n . Trivially, $\mathbf{x}_n \rightarrow \mathbf{x}$. Hence, \mathbf{x} is robustly OE and IC at F . \square

REFERENCES

- Abreu, Dilip, David Pearce, and Ennio Stacchetti (1990), "Toward a theory of discounted repeated games with imperfect monitoring." *Econometrica*, 58, 1041–1063. [0259]
- Agarwal, Nikhil, Itai Ashlagi, Eduardo Azevedo, Clayton R. Featherstone, and Ömer Karaduman (2019), "Market failure in kidney exchange." *American Economic Review*, 109, 4026–4070. [0256]
- Akbarpour, Mohammad, Shengwu Li, and Shayan Oveis Gharan (2020), "Thickness and information in dynamic matching markets." *Journal of Political Economy*, 128, 783–815. [0260]
- Anderson, Ross, Itai Ashlagi, David Gamarnik, and Yash Kanoria (2017), "Efficient dynamic barter exchange." *Operations Research*, 65, 1446–1459. [0260]

⁶³This comes from Prokhorov's theorem and the observation that probability measures in our space are all supported on the same compact set and so the space is automatically tight.

- Arnosti, Nick and Peng Shi (2020), “Design of lotteries and wait-lists for affordable housing allocation.” *Management Science*, 66, 2291–2307. [0283]
- Ashlagi, Itai, Maximilien Burq, Patrick Jaillet, and Vahideh Manshadi (2019), “On matching and thickness in heterogeneous dynamic markets.” *Operations Research*, 67, 927–949. [0260]
- Ashlagi, Itai and Peng Shi (2016), “Optimal allocation without money: An engineering approach.” *Management Science*, 62, 1078–1097. [0257, 0258, 0259, 0264, 0265, 0268, 0269, 0271, 0272, 0274, 0280, 0281, 0288, 0289, 0291]
- Azevedo, Eduardo M. and Jacob D. Leshno (2016), “A supply and demand framework for two-sided matching markets.” *Journal of Political Economy*, 124, 1235–1268. [0283]
- Baccara, Mariagiovanna, SangMok Lee, and Leeat Yariv (2020), “Optimal dynamic matching.” *Theoretical Economics*, 15, 1221–1278. [0260]
- Balseiro, Santiago R., Huseyin Gurkan, and Peng Sun (2019), “Multiagent mechanism design without money.” *Operations Research*, 67, 1417–1436. [0259]
- Benade, Gerdus, Aleksandr Kazachkov, Ariel Procaccia, and Christos-Alexandros Psomas (2018), “How to make envy vanish over time.” In *Proceedings of the 2018 ACM Conference on Economics and Computation EC2018*, 593–610. [0260]
- Bergemann, Dirk and Maher Said (2011), “Dynamic auctions.” *Wiley Encyclopedia of Operations Research and Management Science*. [0259]
- Bikhchandani, Sushil and John W. Mamer (1997), “Competitive equilibrium in an exchange economy with indivisibilities.” *Journal of Economic Theory*, 74, 385–413. [0260]
- Bikhchandani, Sushil and Joseph M. Ostroy (2002), “The package assignment model.” *Journal of Economic Theory*, 107, 377–406. [0260]
- Bloch, Francis and Nicolas Houy (2012), “Optimal assignment of durable objects to successive agents.” *Economic Theory*, 51, 13–33. [0260]
- Bogomolnaia, Anna and Hervé Moulin (2001), “A new solution to the random assignment problem.” *Journal of Economic Theory*, 100, 295–328. [0258, 0283]
- Bogomolnaia, Anna, Hervé Moulin, and Fedor Sandomirskiy (2022), “On the fair division of a random object.” *Management Science*, 68, 1174–1194. [0260]
- Budish, Eric (2011), “The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes.” *Journal of Political Economy*, 119, 1061–1103. [0256, 0258, 0283]
- Budish, Eric, Gerard Cachon, Judd Kessler, and Abraham Othman (2017), “Course match: A large-scale implementation of approximate competitive equilibrium from equal incomes for combinatorial allocation.” *Operations Research*, 65, 314–336. [0256, 0258, 0280]
- Che, Yeon-Koo and Fuhito Kojima (2010), “Asymptotic equivalence of probabilistic serial and random priority mechanisms.” *Econometrica*, 78, 1625–1672. [0258, 0283]

Che, Yeon-Koo and Olivier Tercieux (2019), “Efficiency and stability in large matching markets.” *Journal of Political Economy*, 127, 2301–2342. [0273]

Combe, Julien, Olivier Tercieux, and Camille Terrier (2022), “The design of teacher assignment: Theory and evidence.” *The Review of Economic Studies*, 89, 3154–3222. [0256]

Fudenberg, Drew, David Levine, and Eric Maskin (1994), “The folk theorem with imperfect public information.” *Econometrica*, 62, 997–1039. [0259]

Gorokh, Artur, Siddhartha Banerjee, and Krishnamurthy Iyer (2017), “From monetary to non-monetary mechanism design via artificial currencies.” In *Proceedings of the 2017 ACM Conference on Economics and Computation*, 563–564. [0259]

Guo, Mingyu, Vincent Conitzer, and Daniel M. Reeves (2009), “Competitive repeated allocation without payments.” In *International Workshop on Internet and Network Economics*, 244–255, Springer. [0259]

He, Yinghua, Antonio Miralles, Marek Pycia, and Jianye Yan (2018), “A pseudo-market approach to allocation with priorities.” *American Economic Journal: Microeconomics*, 10, 272–314. [0256]

Hylland, Aanund and Richard Zeckhauser (1979), “The efficient allocation of individuals to positions.” *Journal of Political Economy*, 87, 293–314. [0256, 0258]

Jackson, Matthew O. and Hugo F. Sonnenschein (2007), “Overcoming incentive constraints by linking decisions 1.” *Econometrica*, 75, 241–257. [0259, 0260]

Kamada, Yuichiro and Fuhito Kojima (2023), “Fair matching under constraints: Theory and applications.” *The Review of Economic Studies*, 91, 1162–1199. [0280]

Karp, Richard Maning, Umesh Virkumar Vazirani, and Vijay Vazirani (1990), “An optimal algorithm for on-line bipartite matching.” In *Proceedings of the Twenty-Second Annual ACM Symposium on Theory of Computing (STOC)*, 352–358. [0260]

Kelso, Alexander and Vincent Crawford (1982), “Job matching, coalition formation, and Gross substitutes.” *Econometrica*, 50, 1483–1504. [0260]

Kurino, Morimitsu (2014), “House allocation with overlapping generations.” *American Economic Journal: Microeconomics*, 6, 258–289. [0260]

Lee, SangMok (2016), “Incentive compatibility of large centralized matching markets.” *The Review of Economic Studies*, 84, 444–463. [0273]

Liu, Qingmin and Marek Pycia (2016), “Ordinal efficiency, fairness, and incentives in large markets.” Report. [0258]

Miralles, Antonio and Marek Pycia (2020), “Foundations of pseudomarkets: Walrasian equilibria for discrete resources.” Technical report, working paper. [0258, 0284]

Nguyen, Thành, Ahmad Peivandi, and Rakesh Vohra (2016), “Assignment problems with complementarities.” *Journal of Economic Theory*, 165, 209–241. [0264, 0283]

Østerdal, Lars Peter (2010), “The mass transfer approach to multivariate discrete first order stochastic dominance: Direct proof and implications.” *Journal of Mathematical Economics*, 46, 1222–1228. [0286]

Schrijver, Alexander (1986), *Theory of Linear and Integer Programming*. Wiley, New York. [0269]

Shi, Peng (2022), “Optimal priority-based allocation mechanisms.” *Management Science*, 68, 171–188. [0259]

Zeng, David and Alexandros Psomas (2020), “Fairness-efficiency tradeoffs in dynamic fair division.” In *Proceedings of the 21st ACM Conference on Economics and Computation EC2020*, 911–912. [0260]

Co-editor Bruno Strulovici handled this manuscript.

Manuscript received 12 September, 2022; final version accepted 28 May, 2024; available online 17 July, 2024.