

The Vigilant Eating Rule: A General Approach for Probabilistic Economic Design with Constraints

Haris Aziz[†]Florian Brandl[‡]

We consider probabilistic allocation of objects under ordinal preferences and constraints on allocations. We devise an allocation mechanism, called the vigilant eating rule (VER), that applies to nearly arbitrary constraints. It is constrained ordinally efficient, can be computed efficiently for a large class of constraints, and treats agents equally if they have the same preferences and are subject to the same constraints. When the set of feasible allocations is convex, it is characterized by ordinal egalitarianism. As a case study, we assume objects have priorities for agents and apply VER to sets of probabilistic allocations that are constrained by stability. While VER always returns a stable and constrained efficient allocation, it fails to be strategyproof, unconstrained efficient, and envy-free. We show, however, that each of these three properties is incompatible with stability and constrained efficiency.

1. Introduction

The theory and application of allocation and matching mechanisms have proved to be one of the major success stories of algorithmic economics. An ongoing challenge is designing mechanisms that can handle complex constraints arising in new applications. In this paper, we present a versatile and robust allocation rule that achieves fair and efficient outcomes for a plethora of economic design problems including matching market design. It simultaneously generalizes celebrated rules in the literature including the probabilistic serial rule and the deferred acceptance algorithm.

We study the problem of allocating a set of indivisible objects to a group of agents based on the agents' preferences over objects. The objects could be seats at schools, dormitory rooms, job placements, or organ transplants for example. In these applications, it is important, perhaps even mandatory, that allocations are fair in the sense that no agent (justifiably) envies some

[†]UNSW Sydney, Australia, haris.aziz@unsw.edu.au

[‡]University of Bonn, Germany, florian.brandl@uni-bonn.de

other agent. Since objects are indivisible, envy-free allocations may however not exist. One possible remedy for this dilemma is to consider randomizations over deterministic allocations, which can restore fairness ex-ante. This approach is deeply rooted in the literature on resource allocation and has gained popularity in recent years. The seminal works of Hylland and Zeckhauser (1979) and Bogomolnaia and Moulin (2001) provide allocation mechanisms for the case when each agent demands exactly one object and preferences are given by linear utility functions or strict ordinal preferences over objects, respectively. We shall be concerned with ordinal preferences under three generalizations: first, agents may receive more than one object, second, there may be constraints on which probabilistic allocations are feasible, and third, agents may be indifferent between objects.

The *probabilistic serial rule* of Bogomolnaia and Moulin (2001) can be pictured as follows. Time runs continuously from 0 onward. At time 0, each agent starts off by eating her most-preferred object. An object becomes unavailable if the cumulative time agents spent eating it equals 1. Whenever an object becomes unavailable, the agents who have been eating it switch to their next most-preferred object among those which are still available. The probability with which an agent receives an object in the final probabilistic allocation equals the time she spent eating that object. The Birkhoff-von Neumann Theorem ensures that these probabilities can be attained by randomizing over deterministic allocations. The probabilistic serial rule enjoys several appealing properties: it is weakly strategyproof, can be computed efficiently, and always yields an allocation that is efficient and envy-free (when the agents' preferences over probabilistic allocations are based on stochastic dominance).¹

A tacit assumption in the formulation of the probabilistic serial rule is that all allocations are feasible. But several applications of object allocation require allocation rules to respect various feasibility constraints. For example, allocating students to courses may be subject to curricular constraints that require students to take a minimal number of courses or courses in different subjects. Likewise, when allocating donor's transplant organs to patients, logistics and blood type compatibility impose constraints on which outcomes are feasible. In these examples, one has constraints on deterministic allocations and stipulates that a probabilistic allocation is feasible if it can be written as a convex combination of deterministic allocations that meet the constraints. We will allow for a more general class of constraints where the primitive is a set of feasible *probabilistic* allocations, which includes the previous examples.² More generally, it can capture ex-ante constraints on the probabilistic allocation and possibly different constraints on the ex-post allocation. For example, a system designer may impose ex-ante constraints in addition to ex-post constraints. Then a probabilistic allocation is feasible if it satisfies the ex-ante constraints and can be decomposed into deterministic allocations that meet the ex-post

¹Weak strategyproofness (for preferences based on stochastic dominance) requires that it is never possible for an agent to report her preferences non-truthfully and thereby obtain an allotment that stochastically dominates the allotment obtained by reporting her preferences truthfully.

²In this case, the set of feasible probabilistic allocations is the face of the simplex of probabilistic allocations spanned by deterministic allocations satisfying the constraints.

constraints (see, for example, [Ashlagi et al., 2019b](#); [Akbarpour and Nikzad, 2020](#)).

Our main contribution is a generalization of the probabilistic serial rule, called the *vigilant eating rule (VER)*, which can handle multi-object allocation under nearly arbitrary constraints. The high-level idea is as follows. Each agent increases the probability with which she is guaranteed to receive an object from some indifference class induced by her preferences. Whenever she cannot increase this guarantee further without conflicting with the feasibility constraints or the other agents’ guarantees, she moves on to her next most-preferred indifference class. Formally, VER requires that the set of feasible probabilistic allocations is compact (closed and bounded). Previous generalizations of the probabilistic serial rule required that feasible probabilistic allocations are given by constraints that form a bi-hierarchical structure ([Budish et al., 2013](#)), constitute lower and upper quotas ([Ashlagi et al., 2019b, 2020](#)), or exclude a fixed set of deterministic allocations ([Aziz and Stursberg, 2014](#)). In each of these three domains, VER coincides with the proposed generalization of the probabilistic serial rule. Despite being generally applicable, VER retains many of the properties of the probabilistic serial rule: it always yields an allocation that is efficient among feasible allocations, gives the ex-ante same allotment to agents who have the same preferences and are subject to the same constraints (when a convexity condition on feasible allocations is met), and can be computed efficiently whenever the set of feasible allocations is a union of polytopes (described by polynomially many linear constraints). Clearly, no rule can be unconstrained efficient or envy-free without restrictions on the set of feasible allocations.

A class of problems that falls under the umbrella of object allocation is two-sided allocation, where, in addition to preferences, agents have different priorities for objects. The literature on two-sided allocation is well-developed (see, for example, [Roth and Sotomayor, 1990](#)). Almost universally, allocations are required to be stable: if an agent prefers an object to the one she has been allocated, the former object is given to an agent who has a higher priority for it. We will view stable two-sided allocation as a special case of allocation under constraints, where only stable allocations are feasible. Unlike when constraints come from quotas, stability constraints depend on the preferences of the agents and are thus harder to handle.

In the classical formulation of two-sided allocation as marriage markets, preferences and priorities are strict and one considers only deterministic mechanisms. The most prominent representative is the (agent proposing) deferred acceptance rule of [Gale and Shapley \(1962\)](#). It enjoys several appealing properties such as strategyproofness and Pareto efficiency (when interpreting priorities as preferences). Yet, applications like school choice where priorities are typically coarse, that is, contain large indifference classes, motivate the search for rules that can deal with this more general problem domain. There are various reasons for considering probabilistic rules when priorities are weak. Perhaps the most evident one is that it allows for fairness ex-ante, which is not always attainable when deterministically allocating objects. Take the example of two agents preferring object a over object b . If both agents have the same priority for a and for b , allocating a to either agent with a probability of 50% is fair ex-ante,

while no deterministic allocation is fair.

We propose VER as a promising probabilistic mechanism for stable two-sided allocation when preferences and priorities are weak. Stability can be generalized to probabilistic allocations in various ways. We consider two stability notions—ex-ante stability and ex-post-stability—, which reduce to pairwise stability for deterministic allocations. For both stability notions, we can apply VER to the set of stable allocations and thereby obtain a mechanism that always yields stable and constrained efficient allocations. These mechanisms provide natural transitions between the deferred acceptance algorithm and the probabilistic serial rule. They coincide with the former if priorities are strict and with the latter if priorities are flat (and preferences are strict in both cases).

Contributions Our main contribution is the *Vigilant Eating Rule (VER)*, which applies to any class of distributional constraints, including ones given by stability notions, and allows agents to have weak preferences. For example, it applies to *non-bi-hierarchical* constraints, which are not handled by the generalization of the probabilistic serial rule due to [Budish et al. \(2013\)](#). It also applies to non-convex constraints arising from integrality requirements or some stability notions. Several stability concepts are not necessarily convex. Examples include ex-ante stability ([Kesten and Unver, 2015](#)), respect of priorities ([Delacrétaz, 2020](#)), and a fractional stability concept considered by [Caragiannis et al. \(2019\)](#). We show that the outcome of VER is efficient within the set of feasible allocations with respect to stochastic dominance (that is, *SD*-efficient). If the feasible set is convex, the outcome is ordinally egalitarian in the sense defined by [Bogomolnaia \(2015\)](#) (see also [Bogomolnaia, 2017](#)). Our results about VER are not restricted to allocation problems. They apply to collective decision making in general whenever agents have ordinal preferences over deterministic outcomes.

For object allocation under priorities, we study VER in more detail and argue that it leads to compelling allocation rules. VER returns a stable allocation that is constrained *SD*-efficient within the set of stable allocations. To the best of our knowledge, our algorithms are the first rules for constrained *SD*-efficient probabilistic allocation under priorities when the agents' preferences are weak.³ For several properties that VER does not satisfy, we prove that they are incompatible with stability, constrained (within the set of stable allocations) *SD*-efficiency, or their conjunction. In particular, even for the weaker stability notion we consider (ex-post stability), the set of stable and constrained *SD*-efficient allocations can be disjoint from the set of *SD*-efficient allocations and the set of weakly *SD*-envy-free allocations. For both of our stability notions, no rule can be weakly *SD*-strategyproof and always yield a stable and constrained *SD*-efficient allocation.

More broadly, we present a generalization of the probabilistic serial rule that can handle a more general class of constraints than previous generalizations. In restricted domains, VER

³Some problems invariably involve ties in the agents' preferences. For example, when allocating scarce medical resources such as ventilators where different categories of resources have priorities over agents, agents are indifferent between all ventilator units ([Pathak et al., 2020](#)).

| Constraints | Preferences | Priorities | Resulting Rule |
|--|---------------------|-------------|---|
| – | strict | – | probabilistic serial (PS) (Bogomolnaia and Moulin, 2001) |
| – | weak | – | extended PS (Katta and Sethuraman, 2006) |
| <i>SD</i> -individual rationality ⁴ | strict | – | controlled consuming (Athanassoglou and Sethuraman, 2011) |
| upper quotas ⁵ | strict | – | generalized PS (Budish et al., 2013) |
| lower & upper quotas | strict | – | generalized PS (Ashlagi et al., 2020) |
| feasible ex-post allocations | weak | – | egalitarian simultaneous reservation (Aziz and Stursberg, 2014) |
| feasible ex-post allocations | strict | – | generalized constrained PS (Balbuzanov, 2019) |
| ex-post stability | strict | strict | agent proposing deferred acceptance (DA) (Gale and Shapley, 1962) |
| claimwise stability | strict | weak | constrained PS (Afacan, 2018) |
| – | dichotomous | dichotomous | egalitarian rule (Bogomolnaia and Moulin, 2004) |
| integral allocations; unit demand | weak | – | serial dictatorship for weak preferences (Svensson, 1994) |
| – | strict over bundles | – | PS for bundles (Chatterji and Liu, 2020) |

Table 1: Overview of generalizations of the probabilistic serial rule to particular constraints and preference and priority structures. VER coincides with each of these rules in the corresponding domain.

coincides with well-established rules (see Table 1). We also show that VER can also be useful when no randomization is allowed.

As we endeavor to apply the theory of fair allocation and market design to facilitate applications, it is critical to develop robust mechanisms that can flexibly handle domain-specific constraints. Our approach is especially well-suited for settings with ordinal preferences. We hope that the framework, methodology, and rule we present can be useful for new application domains.

2. Related Work

We explore two-sided allocation while relaxing the assumption that allocations have to be integral or deterministic. In our paper, an allocation specifies, for each pair of agent and object, the probability that the agent receives that object. This probabilistic approach is important on several accounts. We refer the reader to Aziz (2019) for an in-depth discussion. In short, randomization allows one to achieve minimal fairness requirements such as equal treatment of equals that are unobtainable through deterministic allocations. Secondly, probabilistic allocations can also be interpreted as time-sharing or fractional allocations.

Probabilistic matching market design Although most of the work on two-sided allocation has focused on deterministic allocations, some authors have pointed out connections between

⁴The *SD*-individual rationality constraints require that each agent’s allocation stochastically dominates her initial endowment. The endowments are part of the specification of an allocation instance in addition to the agents’ preferences.

⁵Budish et al. (2013) assume that the constraints can be written as the union of two hierarchies of constraints. A collection of constraints forms a hierarchy if for any two constraints, the set of agent-object pairs they apply to are either disjoint or one is contained in the other.

integral stable allocations and their linear relaxations (see, for example, Roth et al. (1993); Teo and Sethuraman (1998)). Erdil and Ergin (2008, 2017) and Kesten (2010) undertook a study of school choice when schools have weak priorities. They highlighted that tie-breaking can lead to a loss of efficiency. However, they focus on achieving constrained efficiency of deterministic or integral allocations. Similarly, Ashlagi et al. (2019a) consider the impact of various types of tie-breaking on the efficiency of the allocations. Kamada and Kojima (2015, 2017a,b) and Kojima et al. (2018) consider deterministic two-sided matching under general distributional constraints. These constraints capture several real-life scenarios, such as restrictions on the number of doctors in a particular region. As noted before, deterministic outcomes cannot guarantee ex-ante equal treatment of equals or capture time-sharing. Deterministic allocation also has several other aspects that are different from its probabilistic counterpart. For example, the problem of computing deterministic allocations that satisfy a set of distributional constraints is NP-hard for many classes of constraints (see, for example, Biró et al., 2010). Considering a probabilistic allocation and then suitably rounding it can be an indirect, but computationally more tractable approach to arrive at a deterministic allocation (see, for example, Akbarpour and Nikzad (2020)).

Kesten and Unver (2015) initiated a serious study of stability notions and mechanisms for probabilistic allocation under weak priorities. They focused on a strong version of ex-ante stability and proposed two mechanisms. First, they present the *fractional deferred acceptance* algorithm, which returns a strongly ex-ante stable allocation. They then modify it to derive the *fractional deferred acceptance and trading* algorithm that returns a strongly ex-ante stable allocation which is also Pareto efficient with respect to stochastic dominance among stable allocations. Han (2017) presented a similar algorithm that returns an ex-ante stable allocation.

Afacan (2018) considered a more general model in which objects have probabilities for prioritizing one agent over another. He proposed a weak stability notion called *claimwise stability* and showed that no weakly strategyproof mechanism always returns an allocation that is claimwise stable and efficient among claimwise stable allocations. Both strategyproofness and efficiency assume that preferences are based on stochastic dominance. Afacan presented his *constrained probabilistic serial* algorithm, which returns an allocation that is constrained efficient among claimwise stable allocations and can be computed in polynomial time.

All of the algorithms that return constrained efficient and stable probabilistic allocations discussed above assume that agents have strict preferences (Afacan, 2018; Han, 2017; Kesten and Unver, 2015). We will not make this assumption and our main results are valid for weak preferences. Furthermore, our approach is general enough to allow for any type of feasibility constraints.

When considering probabilistic stable allocations, there are various notions of stability that coincide with the classical pairwise stability notion for deterministic allocations. Aziz and Klaus (2019) investigate the taxonomy of stability notions for probabilistic allocations and map out the logical relationships between them.

Caragiannis et al. (2019) considered fractional stable allocations under cardinal utilities. A fractional allocation is stable if no pair gets higher utility in being integrally matched to each other than under the fractional allocation. They explore computational problems of finding stable fractional allocations with high welfare. He et al. (2018) also considered cardinal utilities and presented a pseudo-market mechanism that simultaneously generalizes the deferred acceptance algorithm and the mechanism of Hylland and Zeckhauser (1979). It returns an allocation that is ex-ante stable. Echenique et al. (2019) have extended the pseudo-market rule to handle constraints on deterministic allocations.

Extensions of probabilistic serial without priorities There are several papers on probabilistic allocation when the objects have no priorities. Our model allows for weak priorities and, thus, has no or flat priorities as a special case. One of the seminal works on probabilistic allocation under ordinal preferences without priorities is by Bogomolnaia and Moulin (2001), who compared the random priority rule with the probabilistic serial rule. They showed that the probabilistic serial rule is SD -efficient and SD -envy-free, whereas the random priority rule satisfies neither property.

The attractive properties of the probabilistic serial rule have led to a whole line of work on extensions to more general settings. Those include weak preferences (Katta and Sethuraman, 2006), multi-unit demands (Kojima, 2009), endowments (Athanasoglou and Sethuraman, 2011; Yilmaz, 2010), bundled allocations (Chatterji and Liu, 2020), and the probabilistic voting setting (Aziz and Stursberg, 2014). VER coincides with each of these rules on their domain. It captures and generalizes the key insights that underlie the probabilistic serial rule. Moreover, our rule provides a conceptually easier formulation of the previously presented algorithms for specific domains (Katta and Sethuraman, 2006; Yilmaz, 2010; Athanasoglou and Sethuraman, 2011).

Budish et al. (2013) consider probabilistic allocation without priorities. They present a generalization of the probabilistic serial rule which allows for distributional constraints that are given by upper quotas and form a bi-hierarchy. Fujishige et al. (2018) extend the probabilistic serial rule to submodular constraints. In their algorithms, agents continue eating a most preferred object until some upper quota on a set of agent-object pairs is met. This approach works if there are upper quotas on agent-object groups that form a bi-hierarchical constraint structure; it does however not extend to lower quotas on agent-object groups or constraint structures that do not form a bi-hierarchy. Ashlagi et al. (2019b, 2020) consider probabilistic allocation with lower and upper bounds on the types of students for each school without priorities and give a suitable generalization of the probabilistic serial rule. The class of constraints to which VER applies includes those considered by Budish et al. (2013), Fujishige et al. (2018), and Ashlagi et al. (2019b, 2020). In particular, VER can efficiently handle arbitrary linear constraints. There are several differences between VER and the algorithm of Ashlagi et al. (2019b, 2020), including the scope of input preferences, the type of constraints handled, and the type of outputs. The upshot is novel features in the algorithm. VER can handle (i) weak preferences

(instead of strict preferences), (ii) ordinal preferences over bundles of items (rather than single items), (iii) items with priorities over agents (rather than just agents having preferences), (iv) non-linear and non-convex constraints (rather than just lower and upper quotas). Our results include those cases and are thus not implied by the results of [Ashlagi et al. \(2019b, 2020\)](#). This requires methodological novelties, which make our approach different from that of [Ashlagi et al. \(2019b, 2020\)](#). To deal with indifferences in the agents' preferences, rather than eating single objects, agents increase their probability guarantee for indifference classes of objects. To deal with non-convex feasible sets, our rule is based on a priority order over agents. An agent can only eat more of an object if this does not prevent the agents preceding her in the order from eating more. While agents earlier in the order have priority, no requirement is put on how much they can eat except that it is a non-zero amount. When the set of feasible allocations is convex, the outcome is independent of the priority order over agents.

Our model allows for constraints both on the set of feasible ex-post allocations as well as constraints on the ex-ante allocation. Note that if we are only concerned about constraints on the set of ex-post allocations, we can explicitly list the feasible ex-post allocations as the set of alternatives, and run the egalitarian simultaneous reservation rule of [Aziz and Stursberg \(2014\)](#), which is constrained *SD*-efficient and has been referred to as the appropriate extension of the probabilistic serial rule to the social choice domain ([Bogomolnaia, 2017](#)).⁶ However, in allocation problems, there are several advantages of VER. Firstly, the constraints themselves may be ex-ante. For example, three of the four stability notions we consider are defined by constraints on the ex-ante allocation matrix and not by enumerating the set of feasible ex-post allocations. Another example is *time-sharing* where the focus is on the fractional allocations and not on lotteries over deterministic allocations. Secondly, even if certain ex-ante constraints can be captured by enumerating ex-post allocations, the approach may be computationally prohibitive as one may need to enumerate an exponential number of allocations.

3. The Model

Let $N = \{1, \dots, n\}$ be a set of agents and O be a set of m objects. Each agent $i \in N$ has a weak preference relation (complete, reflexive, and transitive) \succsim_i over objects. The symmetric and asymmetric parts of \succsim_i are denoted by \sim_i and \succ_i . A preference relation is strict if it is a linear order; otherwise it is weak. By \mathcal{E}_i we denote the collection of equivalence classes of \sim_i . Abusing notation, we also write \succsim_i for the relation on \mathcal{E}_i induced by \succsim_i . A preference profile $\succsim_N = (\succsim_i)_{i \in N}$ consists of preferences for each agent.

A probabilistic allocation of objects to agents is a matrix $p \in \mathbb{R}_+^{N \times O}$. For $i \in N$ and $O' \subset O$,

⁶More recently, [Balbuzanov \(2019\)](#) proposed an adaptation of the probabilistic serial rule for strict preferences, called the *generalized constrained probabilistic serial rule*, where the set of feasible ex-post allocations are enumerated. Algorithmically, the results of [Balbuzanov \(2019\)](#) can be achieved by enumerating the feasible ex-post allocations and running the egalitarian simultaneous reservation algorithm of [Aziz and Stursberg \(2014\)](#), which seamlessly handles weak preferences.

we write $p(i, O') = \sum_{o \in O'} p(i, o)$. Likewise, for $o \in O$ and $N' \subset N$, we write $p(N', o) = \sum_{i \in N'} p(i, o)$. The allotment of agent i is the row of p corresponding to agent i , denoted by $p(i)$. Agents may receive multiple objects and multiple units of each object may be available. If $p \in \mathbb{N}^{N \times O}$, it is a deterministic allocation. Every probabilistic allocation can be written as a lottery over deterministic allocations.⁷ An allocation p satisfies equal treatment of equals for a profile \succsim_N if for all $i, j \in N$, $\succsim_i = \succsim_j$ implies $p(i, E) = p(j, E)$ for all $E \in \mathcal{E}_i = \mathcal{E}_j$.

We introduce two extensions of preferences over objects to preferences over allotments. Let $x, y \in \mathbb{R}_+^O$. We say that x (first-order) stochastically dominates y with respect to \succsim_i , written $x \succsim_i^{SD} y$, if

$$\sum_{o' \succsim_i o} (x(o') - y(o')) \geq 0 \text{ for all } o \in O \quad (\text{stochastic dominance})$$

If at least one inequality is strict, we write $x \succ_i^{SD} y$. Secondly, x lexicographically dominates y with respect to \succsim_i , written $x \succ_i^{LD} y$, if

$$\text{there is } E \in \mathcal{E}_i \text{ with } x(E) > y(E) \text{ and } x(E') = y(E') \text{ for all } E' \in \mathcal{E}_i \text{ with } E' \succ_i E \quad (\text{lexicographic dominance})$$

If $x(E) = y(E)$ for all $E \in \mathcal{E}_i$, we have $x \sim_i^{LD} y$ and $x \sim_i^{SD} y$ and simply write $x \sim_i y$. Notice that stochastic dominance implies lexicographic dominance, and lexicographic dominance refines stochastic dominance.

We assume that each agent's preferences over allocations only depend on her own allotment. That is, for allocations p and q , $p \succsim_i^{SD} q$ if and only if $p(i) \succsim_i^{SD} q(i)$. Moreover, p stochastically dominates q if $p(i) \succsim_i^{SD} q(i)$ for all $i \in N$ and at least one agent has a strict preference; p is *SD-efficient* if it is not stochastically dominated by any allocation. When $X \subset \mathbb{R}_+^{N \times O}$ is a set of allocations, we say that $p \in X$ is *constrained SD-efficient* for X if it is not stochastically dominated by any allocation in X . The analogous definitions apply when preferences are extended based on lexicographic dominance. Since preferences based on lexicographic dominance refine preferences based on stochastic dominance, *LD-efficiency* is more restrictive than *SD-efficiency*.

An allocation mechanism f maps a preference profile \succsim_N and a non-empty set of feasible allocations X to an allocation $f(\succsim_N, X) \in X$. For any property for allocations, we say that f has this property if $f(\succsim_N, X)$ has the property for all preference profiles \succsim_N and sets X . For example, f is *SD-efficient* if $f(\succsim_N, X)$ is *SD-efficient* for all \succsim_N and X .

4. Vigilant Eating Rules

The formulation of the probabilistic serial rule makes two assumptions. First, it requires that agents always have a unique favorite object and thus strict preferences. If an agent is indifferent

⁷In fact, one can use the Birkhoff-von Neumann Theorem to show the following: if p is a probabilistic allocation with $p(O) = \alpha_i$ for all i , then p can be written as a lottery over deterministic allocations in which agent i receives either $\lfloor \alpha_i \rfloor$ or $\lceil \alpha_i \rceil$ objects.

between multiple objects, one can think of eating from that equivalence class as increasing the probability of getting some object from that class. Second, it presumes that there is one unit of each object and all allocations that entirely assign each object to agents are feasible. The challenge in defining eating rules when not all such allocations are feasible is to not let agents eat objects if this would make it impossible to continue eating in a way that results in a feasible allocation. In other words, we need to ensure that the allocation at each time can be extended to a feasible allocation. Even if the set of feasible allocations consists of all allocations satisfying upper bounds on the probabilities and allocating one unit of each object, this cannot always be achieved by simply letting agents eat objects as long as this does not violate any of the bounds. To see this, consider the following

Example 1. Assume two agents have the preferences depicted below.

$$\begin{aligned} \succ_1: & \quad a \quad b \\ \succ_2: & \quad a \quad b \end{aligned}$$

Suppose we want to get an allocation p in which each agent gets at most one unit of objects and one unit of each object is allocated: $p(1, a) + p(1, b) \leq 1$ and $p(2, a) + p(2, b) \leq 1$, and $p(1, a) + p(2, a) = 1$ and $p(1, b) + p(2, b) = 1$. Suppose additionally that there is a constraint that $p(1, a) + p(2, b) \leq 1/2$. If agents start eating their most preferred object a , they each get $1/2$ of a . Up to this point, no constraint is violated. However, no feasible allocation satisfies $p(1, a) \geq 1/2$ and $p(2, a) \geq 1/2$. The example shows that the Generalized Probabilistic Serial Rule of Budish et al. (2013) cannot handle the above constraints.⁸

4.1. Vigilant Eating

We thus need a more “vigilant” approach. The idea is to let an agent eat an object only if the resulting allocation can be extended to a feasible allocation. In general, it could be however, that agents i and j both can individually eat object o , but not simultaneously. Thus, which objects an agent can eat may depend on which objects other agents eat (and these dependencies may be cyclic). We address this by considering agents one after another according to some priority order. Algorithm 1 formalizes our approach when the priority order is the usual order on N . It allows for weak preferences, so instead of eating objects, agents increase their probability guarantees for objects in the equivalence classes induced by their preferences. That is, the lower bounds on the agents’ probabilities for equivalence classes are increased, and only allocations satisfying those bounds can result in the end. A guarantee can only increase if some feasible allocation satisfying all already established guarantees meets it.

Start with setting all guarantees to 0. At the beginning of each round, we decide for every agent for which equivalence class she gets to increase her guarantee. An equivalence class E is available to agent 1 if there exists a feasible allocation that meets all of the previous guarantees

⁸Budish et al. (2013) defined the generalized probabilistic serial rule for bi-hierarchical constraint structures. The constraints in Example 1 do not form a bi-hierarchy.

and assigns a strictly higher probability for E to i than her current guarantee for E . Agent 1 gets to increase her probability for her most preferred equivalence class E_1 among those available to her. In general, E is available to agent i if there exists a feasible allocation that meets all guarantees established in the previous rounds, assigns a higher probability for E_j to j for each $j < i$, and assigns a higher probability for E to i than the current guarantee. Agent i gets to eat her favorite equivalence class E_i among those available to her. Agents can increase their probabilities for their favorite classes E_i for as long as all the resulting guarantees can be met by some feasible allocation. More precisely, if $\pi^k: \{(i, E): i \in N, E \in \mathcal{E}_i\} \rightarrow \mathbb{R}_+$ stores the guarantees at the beginning of round k , find the largest δ so that there exists a feasible allocation which meets all guarantees imposed by π^k and assigns a probability of at least $\pi^k(i, E_i) + \delta$ for E_i to i for all i . If δ^* is this maximal value, we let π^{k+1} be array that increases i 's guarantee for E_i by δ^* compared to π^k and is otherwise identical to π^k . Assuming that X is compact ensures that δ^* exists. After at most $m \cdot n$ rounds, no guarantee can be further increased and the process terminates. By construction, X contains an allocation that meets all of the guarantees and all agents are indifferent between all such allocations. Hence, it is irrelevant which of those we choose as the final allocation.

We call this mechanism the *vigilant eating rule (VER)*. Whenever the set of feasible allocations X is non-empty and compact, VER results in a feasible allocation that is constrained *SD*-efficient with respect to X (see Theorems 1 and 2). For feasible sets that are not compact, constrained *SD*-efficient allocations may not exist, which shows that our requirement is minimal. If X is a union of polytopes, it can be computed in polynomial time in number of inequalities used to describe X (Theorem 3). Moreover, if X is convex, the outcome of VER is ordinally egalitarian (Theorem 4) and independent of the priority order over agents. If X meets certain symmetry conditions, VER satisfies equal treatment of equals (Theorem 5). Lastly, we explore an extension of VER that allows agents to eat at different rates (Theorem 6).

Example 2 (Illustration of VER). Let us revisit Example 1, where we had two agents with the following preferences.

$$\begin{aligned} \succ_1: & \quad a \quad b \\ \succ_2: & \quad a \quad b \end{aligned}$$

The constraints on the allocation are

$$\begin{aligned} p(1, a) + p(1, b) &\leq 1 & p(2, a) + p(2, b) &\leq 1 \\ p(1, a) + p(2, a) &= 1 & p(1, b) + p(2, b) &= 1 \\ p(1, a) + p(2, b) &\leq 1/2 \end{aligned}$$

We illustrate how VER works in this example. The order of consumption of objects over time is illustrated in Figure 1.

- Round 1: The most preferred object for agent 1 that she can consume is a . Next, the most preferred object for agent 2 that she can consume given that agent 1 receives a with

Algorithm 1 The Vigilant Eating Rule

Input: A preference profile \succsim_N and a non-empty and compact set X of allocations**Output:** An allocation $p \in X$

- 1 $k \leftarrow 0$ (the round of the algorithm)
 - 2 $N' \leftarrow N$ (the set of active agents)
 - 3 $\pi^0(i, E) \leftarrow 0$ for all $i \in N$ and $E \in \mathcal{E}_i$
 - 4 **while** $N' \neq \emptyset$ **do**
 - 5 **for** $i \in N'$ in increasing order **do**
 - 6 $\mathcal{F}_i \leftarrow \{E \in \mathcal{E}_i : \text{there is } p \in X \text{ so that } p(j, E') \geq \pi^k(j, E') \text{ for all } j \in N \text{ and } E' \in \mathcal{E}_j, p(i, E) > \pi^k(i, E), \text{ and } p(j, E_j) > \pi^k(j, E_j) \text{ for all } j \in N' \text{ with } j < i\}$
 - 7 **if** $\mathcal{F}_i \neq \emptyset$ **then**
 - 8 $E_i \leftarrow \max_{\succsim_i} \mathcal{F}_i$
 - 9 **else**
 - 10 $N' \leftarrow N' \setminus \{i\}$
 - 11 **end if**
 - 12 **end for**
 - 13 Compute δ^* as follows

$$\delta^* = \max_{\delta, p} \delta \text{ s.t.}$$

| | | |
|---|---|-----------------------------|
| $p(i, E_i) \geq \pi^k(i, E_i) + \delta$ | $\forall i \in N'$ | (agent i consumes E_i) |
| $p(i, E) \geq \pi^k(i, E)$ | $\forall i \in N$ and $E \in \mathcal{E}_i$ | (p extends π^k) |
| $p \in X$ | | (p is feasible) |
 - 14 $\pi^{k+1}(i, E) \leftarrow \pi^k(i, E)$ for all $i \in N$ and $E \in \mathcal{E}_i$
 - 15 $\pi^{k+1}(i, E_i) \leftarrow \pi^k(i, E_i) + \delta^*$ for all $i \in N'$
 - 16 $k \leftarrow k + 1$
 - 17 **end while**
 - 18 **return** p (as computed in Line 13)
-

| | | | | | |
|---------|-----|---------------|---------------|---------------|---|
| Agent 1 | a | b | b | | |
| Agent 2 | a | a | b | | |
| | 0 | $\frac{1}{4}$ | $\frac{1}{2}$ | $\frac{3}{4}$ | 1 |

Time

Figure 1: Illustration of VER for the problem in Example 1.

positive probability is also a . The algorithm computes the maximum amount that the agents can consume from a while ensuring that there exists a feasible allocation extending the current allocation. This amount is $1/4$ since if agent 1 were to consume more of object a , the column constraints would force agent 2 to consume more than $1/4$ of object b . A violation of the diagonal constraint would thus be inevitable. Hence, agent 1 has to move on to consuming object b . At the end of the round, we have the following guarantees: $\pi^1(1, a) \geq 1/4$ and $\pi^1(2, a) \geq 1/4$.

- Round 2: Agent 1 consumes $1/2$ of object b whereas agent 2 consumes $1/2$ more of object a . At this point, object a becomes unavailable and agent 2 has to move on to consuming object b . At the end of the round, we have the following guarantees: $\pi^2(1, a) \geq 1/4$, $\pi^2(1, b) \geq 1/2$, and $\pi^2(2, a) \geq 3/4$.
- Round 3: Now both agents consume $1/4$ more of b until b becomes unavailable. At the end of round 3, we have the following guarantees: $\pi^3(1, a) \geq 1/4$, $\pi^3(1, b) \geq 3/4$, $\pi^3(2, a) \geq 3/4$ and $\pi^3(2, b) \geq 1/4$. Since both objects are unavailable, the algorithm terminates with the following allocation satisfying all guarantees and the original constraints.

$$p = \begin{matrix} & a & b \\ \begin{matrix} 1 \\ 2 \end{matrix} & \begin{pmatrix} 1/4 & 3/4 \\ 3/4 & 1/4 \end{pmatrix} \end{matrix}$$

Example 3 (Illustration of VER with indifferences). We modify Example 2 so that agent 1 is indifferent between a and b .

$$\succ_1: \{a, b\}$$

$$\succ_2: \quad a \quad b$$

The constraints on the allocation remain the same.

$$\begin{aligned} p(1, a) + p(1, b) &\leq 1 & p(2, a) + p(2, b) &\leq 1 \\ p(1, a) + p(2, a) &= 1 & p(1, b) + p(2, b) &= 1 \\ p(1, a) + p(2, b) &\leq 1/2 \end{aligned}$$

VER proceeds as follows.

- Round 1: The most preferred indifference class for agent 1 that she can consume is $\{a, b\}$. Next, the most preferred object for agent 2 that she can consume given that agent 1

receives some object from $\{a, b\}$ with positive probability is a . The algorithm computes the maximum amount that the agents can consume from $\{a, b\}$ and a , respectively, while ensuring that there exists a feasible allocation. This amount is 1 since the allocation

$$p = \begin{matrix} & a & b \\ \begin{matrix} 1 \\ 2 \end{matrix} & \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \end{matrix} \quad (\star)$$

is feasible. At the end of the round, we have the following guarantees: $\pi^1(1, \{a, b\}) \geq 1$ and $\pi^1(2, a) \geq 1$. Since both objects are unavailable, the algorithm terminates with the allocation p .

Breaking the tie in agent 1's preferences in favor of a and then applying VER gives an allocation assigning a to agent 1 with positive probability (see Example 2). This allocation is stochastically dominated by p as given in (\star) . Hence, handling indifferences by breaking ties in the agents' preferences arbitrarily can produce inferior outcomes.

We first verify that each step of VER is well-defined and that it always terminates. It is then clear from the definition that it results in a feasible allocation. The proof proceeds by showing that agents consume equivalence classes in decreasing order of their preferences and then observing that in each round some agent moves to a less preferred class or becomes inactive since she cannot increase any of her guarantees further. (All proofs are in the appendix.)

Theorem 1. *For every non-empty and compact set of allocations X , Algorithm 1 terminates after at most $m \cdot n$ rounds with an allocation $p \in X$.*

Intuitively, it is clear that VER yields an allocation that is *LD*-efficient (and, thus, *SD*-efficient) among feasible allocations. To see this, assume that the allocation p of VER were dominated by some allocation q in X . We can view q as the result of some eating trajectory different from the one which yields p in Algorithm 1. Assuming agents increase their guarantees for equivalence classes at unit rate in decreasing order of their preference pins down a unique trajectory for q . Consider the earliest point in time at which both trajectories diverge and let i^* be the smallest agent who starts increasing her guarantee for different equivalence classes in both trajectories. Since q dominates p , i^* strictly prefers her equivalence class in the q -trajectory to that in the p -trajectory. But this contradicts the choice of E_{i^*} at that time in Line 8 of Algorithm 1.

Theorem 2. *When VER is applied to any non-empty and compact set of allocations X , it returns an allocation that is constrained *LD*-efficient among the allocations in X .*

The steps in Algorithm 1 that are potentially computationally taxing are determining for which equivalence classes an agent can increase her guarantee in Line 8 and the amount of time the current eating scheme can be continued in Line 13. Both require maximizing a linear

objective over the intersection of X with a polytope. This can be computationally prohibitive if, for example, X consists of an exponential number of deterministic allocations. However, whenever X is itself a polytope (described by polynomially many linear inequalities), this can be done efficiently by solving a linear program. More generally, if X is a union of polytopes, we can solve the problem for each polytope individually and then pick the best solution (which is the most-preferred equivalence class E_i agent i can eat or the maximal amount of time δ^* the current eating scheme can continue).

Theorem 3. *If X is a non-empty finite union of polytopes, the outcome of VER can be computed in polynomial time in the number of inequalities used to describe X .*

4.2. Ordinal Egalitarianism and Equal Treatment of Equals

Bogomolnaia (2015) proposed a solution notion called *ordinal egalitarianism* that characterizes the probabilistic serial rule (which can be argued to be the fairest rule for the allocation of objects to agents under ordinal preferences). Ordinal egalitarianism is defined as follows. We evaluate any given allocation by the list of numbers t_i^j , the total probability agent i gets of objects from her first j equivalence classes, for all i and j . An allocation is *ordinally egalitarian* if it is a leximin maximizer of the signature vector $t^\uparrow(p) = (t_i^j(p))_{i,j}$ over all feasible allocations. For any signature vector $t^\uparrow(p)$, we will denote by $t^\uparrow(p)(j)$ as the j -th entry of the non-decreasing ordered entries. We will denote by $t_i^\uparrow(p)$ as the signature vector $t^\uparrow(p)$ in which only the entries corresponding to agent i are considered. If an allocation is ordinally egalitarian, then it is *LD-efficient*. We show that under convex constraints, VER returns an allocation that is ordinally egalitarian.

Theorem 4. *The VER allocation for a non-empty, compact, and convex set of allocations X is ordinally egalitarian among the allocations in X .*

Remark 1. It follows from Theorem 4 that if X is convex, the outcome of VER is independent of the priority order over agents, which determines in which order agents choose objects. Hence, in that case, like the eating-based algorithms in the literature, VER can be formulated so that it is symmetric with respect to the agents.

One of the desiderata when assigning objects to agents is to treat all agents equally. Roughly, this means that if two agents have the same preferences and are subject to the same constraints, then they should receive the same allocation. Thus, there are two ways in which asymmetry between agents in an allocation can arise: the mechanism that determines it is inherently asymmetric or the set of feasible allocations is biased toward certain agents.

VER introduces some asymmetry of the first kind by letting agents choose their most-preferred available object according to some priority order. In general, this order is relevant since whether or not an equivalence class is available to an agent may depend on the choices of previous agents. It is clear that this asymmetry cannot be eliminated entirely. For example, two agents may have the same most-preferred object and only the allocations assigning that object to one of them

with probability 1 may be feasible. (This is just the special case of deterministic mechanisms.) However, if X is convex, the order in which agents choose equivalence classes is irrelevant (see Remark 1). A second source of asymmetry is the set of feasible allocations X itself. Clearly, if X is asymmetric, we may not be able to give the same allocation to agents with the same preferences.

These considerations motivate the conditions under which VER treats agents equally. Let i and i' be two agents. We say that a set of allocations X is symmetric for $\{i, j\}$ if $p \in X$ implies $p^{(ij)} \in X$, where $p^{(ij)}$ is identical to p except that the allocations of i and j are swapped. Moreover, X is convex for $\{i, i'\}$ if for all $\lambda \in [0, 1]$, $p \in X$ and $p^{(ii')} \in X$ implies $\lambda p + (1 - \lambda)p^{(ii')} \in X$.

Theorem 5. *Let X be a non-empty and compact set of allocations and \succsim_N be a preference profile. If X is symmetric and convex for all $i, i' \in N$ with $\succsim_i = \succsim_{i'}$, then the outcome of VER satisfies equal treatment of equals.*

4.3. Variable Eating Rates

One can generalize VER by equipping each agent i with an eating rate function $f_i: \mathbb{R}_+ \rightarrow \mathbb{R}_+$ (assumed to be measurable), which determines how fast she can increase her probability for an equivalence class as a function of time. Algorithm 1 can be extended to eating rate functions by introducing variables t^k with $t^0 = 0$ and $t^k = t^{k-1} + \delta^*$ as determined in Line 13 in round k , replacing the second line in the linear program in Line 13 by

$$p(i, E_i) \geq \pi^k(i, E_i) + \int_{t_k}^{t_k + \delta} f_i(s) ds \quad \forall i \in N',$$

and replacing Line 15 by

$$\pi^{k+1}(i, E_i) \leftarrow \pi^k(i, E_i) + \int_{t_k}^{t_k + \delta^*} f_i(s) ds \quad \forall i \in N'.$$

If the eating rate functions are piecewise constant, which is the case we will consider, δ^* can still be computed by linear programming. In general, this is not possible, however.

Bogomolnaia and Moulin (2001, Theorem 1) show that for unit demand object allocation without constraints, an allocation is SD -efficient if and only if it is the outcome of their eating mechanism for some profile of eating rate functions. The “if part” remains true in our setting, that is, for every profile of eating rate functions, VER yields a constrained LD -efficient and, thus, constrained SD -efficient allocation. But not every constrained SD -efficient allocation is the outcome of VER for some profile of eating rate functions.

Example 4 (Constrained SD -efficiency). Consider the following preferences.

$$\begin{aligned} \succsim_1: & \quad a \quad b \quad c \\ \succsim_2: & \quad b \quad c \quad a \\ \succsim_3: & \quad c \quad a \quad b \end{aligned}$$

Let X be the convex hull of the allocations p and q .

$$p = \begin{array}{c} \\ \\ \\ \end{array} \begin{array}{ccc} a & b & c \\ \left(\begin{array}{ccc} 1/2 & 0 & 1/2 \\ 1/2 & 1/2 & 0 \\ 0 & 1/2 & 1/2 \end{array} \right) \end{array} \quad q = \begin{array}{c} \\ \\ \\ \end{array} \begin{array}{ccc} a & b & c \\ \left(\begin{array}{ccc} 1/3 & 1/3 & 1/3 \\ 1/3 & 1/3 & 1/3 \\ 1/3 & 1/3 & 1/3 \end{array} \right) \end{array}$$

Every allocation in X is constrained SD -efficient. However, there is no profile of eating rate functions for which VER will yield q . This is because at each point during the eating process, every agent's favorite object o will be available to her until she has consumed $\frac{1}{2}$ of o .

The reason for the equivalence derived by [Bogomolnaia and Moulin](#) is that SD -efficiency and LD -efficiency coincide in their setting as shown by [Cho and Dogan \(2016\)](#). This is no longer true in our case as observed above. We can however show that when the set of feasible allocations is convex, every constrained LD -efficient allocation is unanimously indifferent to the outcome of VER for some profile of eating rate functions. To this end, it suffices to consider indicator functions as eating rate functions. Those result in one-at-a-time-VER, which at any point in time, allows only one agent to increase her probability for some equivalence class. Formally, for all $i \in N$, $f_i = \chi_{S_i}$, where the S_i are pairwise disjoint measurable subsets of \mathbb{R}_+ and χ_S is the indicator function of $S \subset \mathbb{R}_+$. Clearly, the outcome of VER can be achieved by some instance of one-at-a-time-VER.

Theorem 6. *Let X be a non-empty, compact, and convex set of allocations and $p \in X$. Then p is constrained LD -efficient with respect to X if and only if all agents are indifferent between p and the outcome of some instance of one-at-a-time-VER on the set X .*

Remark 2. An interesting instance of one-at-a-time-VER is the *vigilant priority rule*. Fix some order of the agents. The first agent in that order increases her probability guarantees for each of her equivalence classes in order of her preferences by as much as possible. Once no further increase is possible for the first agent, it is the second agent's turn to increase all of her guarantees in order of her preferences by as much as possible. We proceed in the same way for the remaining agents.⁹

The vigilant priority rule is clearly unfair to agents who come later in the order. It is (strongly) SD -strategyproof whenever the set of feasible allocations is independent of the agents' preferences. If we assume the set of feasible allocations is convex and choose the order of agents uniformly at random, the rule is no longer LD -efficient. To see this, observe that it coincides with the classic random priority rule in the assignment domain, which is well-known to violate SD -efficiency (and thus LD -efficiency).

⁹We obtain the vigilant random priority rule as an instance of one-at-a-time VER by setting $S_i = [(l-1)C, lC)$, where l is the rank of agent i in the chosen order and C is larger than the total amount of objects allocated in any feasible allocation.

4.4. Examples of Constraints

We conclude this section by giving concrete classes of constraints that VER can handle.

Distributional constraints on the ex-ante allocations We have pointed out that VER applies to a wide variety of constraints on the ex-ante allocations. A natural and general family of constraints arises from imposing lower and upper bounds on the cumulative probability a subset of agents can obtain from a subset of objects. These constraints appear in applications like school choice if a certain number of seats is reserved for students who live close to the school. Another class of constraints obtains from diversity requirements, which put bounds on the fraction of objects from within some subset that is assigned to a certain subset of agents (instead of bounding the absolute number).

Distributional constraints on the ex-post allocations As mentioned in the introduction, ex-post constraints, that is, constraints on deterministic allocations, can be expressed as ex-ante constraints on the set of probabilistic allocations by taking only those allocations to be feasible which can be decomposed into deterministic allocations satisfying the ex-post constraints. These can capture upper and lower bounds on the number of objects a subset of agents can receive from some subset of objects ex-post, for example. One instance of this are bi-hierarchical constraints on deterministic allocations of the type studied by [Budish et al. \(2013\)](#). Similarly, one can also formulate diversity constraints ex-post.

Ex-post constraints can be combined with possibly different ex-ante constraints on probabilistic allocations. For example, one may stipulate that the ex-ante constraints hold exactly, while the ex-post constraints only need to be satisfied approximately. [Akbarpour and Nikzad \(2020\)](#) show that if the ex-ante and ex-post constraints are the same, then any allocation satisfying the ex-ante constraints can be decomposed into deterministic allocations that approximately achieve the constraints. Hence, in that case, there is little need to impose additional constraints which guarantee that a decomposition into deterministic allocations satisfying the constraints exists.

Allocation with endowments VER can be applied to problems in which agents have endowments to obtain individually rational and possibly more efficient re-allocations. [Athanasoglou and Sethuraman \(2011\)](#) presented the controlled consuming algorithm, which can be viewed as an extension of the probabilistic serial rule.¹⁰ It ensures that each agent gets an allocation she weakly prefers to her endowment with respect to stochastic dominance. The requirement of *SD*-individual rationality can be embedded in our framework by using a linear number of inequalities for each agent. Hence, we can use linear programming to capture the controlled consuming algorithm.

¹⁰The algorithm of [Athanasoglou and Sethuraman \(2011\)](#) is more general than the generalized eating algorithm of [Yilmaz \(2010\)](#) as the latter works only for strict preferences.

Deterministic and integral allocations Our approach has been framed in the context of randomization or time-sharing. However, it also applies when no randomization is involved or allowed by rendering all non-deterministic allocations infeasible. Then the δ increment in the probability of an agent for an object (cf. Algorithm 1, Line 13) always has to be an integer. This ensures that an agent gets some number of units of each object with probability one. Not allowing probabilistic allocations under feasibility constraints may however render many problems NP-hard and thus computationally intractable. We note that for the house allocation problem with deterministic allocations (Svensson, 1994, 1999), VER coincides with the serial dictatorship rule.

5. Stable Probabilistic Allocations

We assume now that agents have different priorities for objects, resulting in a two-sided allocation problem. Requiring allocations to be stable imposes feasibility constraint on allocations. Thus, choosing among stable allocations, possibly with additional properties, is within the framework of the preceding section. Various stability notions for probabilistic allocations have been proposed, of which we will discuss two. We then examine VER as a mechanism for stable probabilistic object allocation under priorities. Many of its properties follow from the preceding general statements. On the other hand, VER fails to be strategyproof, unconstrained efficient, and envy-free. We show that this is unavoidable if one insists on stability.

Assume each object o has a (complete, reflexive, and transitive) relation \succsim_o over agents, which specifies the agents' priorities for o . A profile $(\succsim_N, \succsim_O) = ((\succsim_i)_{i \in N}, (\succsim_o)_{o \in O})$ consists of preferences for each agent and priorities for each object.¹¹ Moreover, each agent i has a capacity $\text{cap}(i) \in \mathbb{N}$ of objects she can receive. Likewise, for each object o , $\text{cap}(o) \in \mathbb{N}$ is the amount of o that is available. We stipulate that for every allocation p , $p(i, O) \leq \text{cap}(i)$ and $p(N, o) \leq \text{cap}(o)$ for all $i \in N$ and $o \in O$. The priorities give rise to a relaxed notion of equal treatment of equals, which only requires that two agents receive the same allotment if they have the same preferences and the same priority for all objects. That is, an allocation p satisfies limited equal treatment of equals for a profile (\succsim_N, \succsim_O) if $p(i) = p(j)$ for all $i, j \in N$ with $\succsim_i = \succsim_j$, $\text{cap}(i) = \text{cap}(j)$, and $i \sim_o j$ for all $o \in O$.

We consider two notions of stability, which coincide with the standard (pairwise) stability for deterministic allocations if the agents' preferences over allocations are responsive. Recall that a deterministic allocation p is stable if for all $i \in N$ and $o \in O$, either $\sum_{o' \succsim_i o} p(i, o') = \text{cap}(i)$ or $\sum_{j \succsim_o i} p(j, o) = \text{cap}(o)$.

Ex-ante stability has been introduced by Kesten and Unver (2015) and is defined by the same constraints as stability for deterministic allocations. It prescribes that for any agent i and object o , either i receives $\text{cap}(i)$ of objects weakly preferred to o or o is allocated entirely to

¹¹Using preferences for agents and priorities for objects follows the school choice literature (Abdulkadiroğlu and Sönmez, 2003), which considers efficiency, fairness, and strategic aspects only from the perspective of the agents.

agents with weakly higher priority for o than i . Formally, p is ex-ante stable if for all $i \in N$ and $o \in O$,

$$\sum_{o' \succsim_i o} p(i, o') = \text{cap}(i) \quad \text{or} \quad \sum_{j \succsim_o i} p(j, o) = \text{cap}(o) \quad (\text{ex-ante stability})$$

Ex-ante stability implies that for all $i, j \in N$ and $o \in O$, $i \succ_o j$ and $p(j, o) > 0$ only if $\sum_{o' \succsim_i o} p(i, o') = \text{cap}(i)$. Thus, ex-ante stability requires that i does not justifiably envy any agent j before knowing the realization of the probabilistic allocation p .

Analogously, one can ask that no agent has justified envy ex-post, that is, after a deterministic allocation has been selected according to the probabilistic allocation. This leads to defining p as ex-post stable if

$$p \text{ is a convex combination of deterministic stable allocations} \quad (\text{ex-post stability})$$

It is straightforward to check that ex-ante stability implies ex-post stability.

5.1. Vigilant Eating on Sets of Stable Allocations

In this section, we study the properties of VER when it is applied to sets of stable allocations. If S is one of our stability notions, then S -VER denotes the mechanism which, for a profile (\succsim_N, \succsim_O) , runs VER for the preferences \succsim_N on the set X of S -stable allocations for (\succsim_N, \succsim_O) , so that S -VER $(\succsim_N, \succsim_O) = \text{VER}(\succsim_N, X)$. Note that X is always non-empty since deterministic stable allocations always exist in our setting, and it is compact because of the capacity constraints. Thus, many of the properties of VER on sets of stable allocations follow from the results in Section 4. In particular, for ex-ante stability and ex-post stability, VER yields an allocation that is stable, constrained SD -efficient, and satisfies limited equal treatment of equals.

Corollary 1. *If S is ex-ante or ex-post stability, S -VER returns an S -stable allocation that is SD -efficient among S -stable allocations and satisfies limited equal treatment of equals.*

For strict preferences and the extreme cases of coarseness of priorities, VER coincides with well-known mechanisms. If priorities are flat (all agents have the same priority for every object), every allocation is stable for both stability notions. In that case, VER reduces to the probabilistic serial rule of [Bogomolnaia and Moulin \(2001\)](#), which corresponds to unconstrained eating. If priorities are strict, VER returns the agent optimal deterministic stable allocation, which is also the outcome of the agent-proposing deferred acceptance algorithm. Intuitively, this checks out since VER is optimal for agents in the sense that it allows them to eat their most preferred object available to them.

Corollary 2. *Assume that preferences and priorities are strict. Then, if S is ex-ante stability or ex-post stability, S -VER returns the agent optimal deterministic stable allocation.*

In view of Corollary 1 and Corollary 2, we also recover the well-known result that the agent proposing deferred acceptance algorithm returns an allocation that is Pareto efficient within the

set of stable allocations when the preferences and priorities are strict. Corollary 2 shows that S -VER results in the same mechanism for both of our stability notions in that case even though the sets of stable allocations are not the same. For weak priorities, the two instantiations of VER are in fact distinct.

An alternative interpretation of our model is that the entities on both sides of the market are agents who have preferences over the other side (instead of one side being objects with priorities over agents). Then, instead of considering SD -efficiency for one side, we could ask for allocations that are SD -efficient with respect to the preferences of both sides, called two-sided SD -efficiency. Formally, an allocation p is *two-sided SD -efficient* for a profile (\succsim_N, \succsim_O) if there is no allocation q such that $q(i) \succsim_i^{SD} p(i)$ for all $i \in N$ and $q(o) \succsim_o^{SD} p(o)$ for all $o \in O$ and at least one preference is strict. We show that VER applied to the set of ex-ante stable allocations always returns an allocation that is two-sided SD -efficient among all allocations, not only among stable allocations, when preferences are strict. For ex-post stability, this is open.

Proposition 1. *Assume that the agents' preferences are strict. If S is ex-ante stability, S -VER returns an allocation that is two-sided SD -efficient.*

5.2. Incompatibility of Stability with Efficiency, Envy-Freeness, and Strategyproofness

For both of our stability notions, VER on the set of stable allocations violates unconstrained efficiency, weak envy-freeness, and weak strategyproofness when preferences over probabilistic allocations are based on stochastic dominance. Neither of those shortcomings is specific to VER, but a consequence of an inherent incompatibility of each of these properties with stability and constrained efficiency. We address them in turn.

Roth (1982) showed that there may be no deterministic allocation that is both stable and Pareto efficient. Since any ex-post stable and SD -efficient allocation has to be a convex combination of stable and Pareto efficient deterministic allocations, it follows that ex-post stability (and, thus, also ex-ante stability) is incompatible with unconstrained SD -efficiency.

In the absence of priorities, envy-freeness requires that each agent prefers her allocation to that of any other agent. This definition is no longer compelling for non-trivial priorities since an agent may legitimately receive an allocation that some other agent would prefer to her own because of a higher priority for some objects. One definition of fairness in this context is limited equal treatment of equals, which we discussed above. Another one is that of justified envy. It applies only to pairs of agents i, j who have the same priority for all objects and requires that i does not prefer j 's allotment to her own when comparing them via stochastic dominance. That is, an allocation p is weakly SD -envy-free if for all $i, j \in N$,

$$i \sim_o j \text{ for all } o \in O \text{ implies } p(j) \not\succeq_i^{SD} p(i) \quad (\text{weak } SD\text{-envy-freeness})$$

Even this weak notion of envy-freeness turns out to be incompatible with ex-post stability and constrained ex-post SD -efficiency.

Proposition 2. *There may be no allocation that is ex-post stable, weakly SD-envy-free, and SD-efficient among ex-post stable allocations (if there are at least three agents and three objects).*

Comparing allocations via stochastic dominance results in incomplete preferences. Thus, there are two notions of strategyproofness associated with it. The stronger, usually called *SD-strategyproofness*, requires that the allotment obtained by truth-telling weakly stochastically dominates any allotment that can be obtained otherwise. [Bogomolnaia and Moulin \(2001\)](#) proved that when priorities are flat, there exists no mechanism that is *SD-strategyproof*, *SD-efficient*, and satisfies equal treatment of equals. Under flat priorities, *SD-efficiency* is the same as constrained *SD-efficiency* since stability has no bite. Thus, there is no mechanism that returns stable allocations and satisfies *SD-strategyproofness*, constrained *SD-efficiency*, and limited equal treatment of equals (irrespective of the stability notion).

The weak notion of *SD-strategyproofness* only prescribes that no agent can obtain a strictly *SD*-dominating allotment by misrepresenting her preferences. Formally, a mechanism f is weakly *SD-strategyproof* if for all agents $i \in N$ and all profiles $(\succsim_N, \succsim_O), (\succsim'_N, \succsim_O)$ with $\succsim_j = \succsim'_j$ for all $j \in N \setminus \{i\}$,

$$f(\succsim'_N, \succsim_O)(i) \not\succeq_i^{SD} f(\succsim_N, \succsim_O)(i).$$

For flat priorities and strict preferences, there is a mechanism that satisfies weak *SD-strategyproofness*, *SD-efficiency*, and equal treatment of equals: the probabilistic serial rule. Moreover, the generalized probabilistic serial rule for upper quotas forming a bi-hierarchy of [Budish et al. \(2013\)](#) is weakly strategyproof. But imposing stability again results in an impossibility. [Afacan \(2018\)](#) proved that for at least 4 agents, there exists no mechanism that is weakly *SD-strategyproof* and always returns a claimwise stable and constrained *SD-efficient* allocation.¹² We prove the analogous result for ex-ante stability and ex-post stability. Note that constrained efficiency becomes weaker as stability becomes stronger, since the set of stable allocations becomes smaller. Thus, none of these statements implies the other. Our proof does not rely on weak preferences and requires only 3 agents.

Proposition 3. *If S is ex-ante stability or ex-post stability, no mechanism satisfies weak SD-strategyproofness, S -stability, and S -constrained efficiency even if preferences are strict (and there are at least three agents and three objects).*

All the properties are needed for the conclusion of Proposition 3. VER satisfies S -stability and S -constrained efficiency. Ignoring the priorities and running the probabilistic serial rule gives a mechanism that is weakly *SD-strategyproof* and efficient (and thus S -constrained efficient), but not S -stable. The deferred acceptance algorithm with lexicographic tie-breaking of priorities is

¹²Claimwise stability has been introduced by [Afacan \(2018\)](#). Say that agent i has a justified claim against j for object o if i has higher priority for o than j and j 's probability for o is larger than i 's probability for objects i weakly prefers to o . An allocation p is claimwise stable if no agent has a justified claim, that is, if for all $i, j \in N$ and $o \in O$, $i \succ_o j$ implies $\sum_{o' \succ_i o, o' \neq o} p(i, o') \geq p(j, o)$.

SD -strategyproof and S -stable. For ex-post stability, one can simultaneously achieve limited equal treatment of equals by breaking ties uniformly at random.

In Appendix C, we give a detailed comparison of various mechanisms for probabilistic object allocation under priorities.

6. Extension Beyond Allocation Problems

The VER framework is not restricted to allocation problems. It applies just as well to other settings where fractional or probabilistic outcomes are feasible, such as social choice (Bogomolnaia et al., 2005; Brandt, 2017), coalition formation (Bogomolnaia and Jackson, 2002; Aziz et al., 2013b), networks (Jackson and Wolinsky, 1996), and other models discussed by Sönmez (1999). In this section, we discuss how to extend our model to capture these applications.

Instead of a set of objects, we now consider an abstract set of alternatives A and each agent has a preference relation over A . An outcome is an element of \mathbb{R}_+^A . The agents' preferences over outcomes are again based on stochastic dominance. By a problem we denote a pair (\succsim_N, X) , where \succsim_N is a preference profile and $X \subset \mathbb{R}_+^A$ is a non-empty set of feasible outcomes. As before, a mechanism f maps a problem (\succsim_N, X) to an outcome $f(\succsim_N, X) \in X$.

A generalization of VER then works as follows: in each round, agents are addressed in lexicographic order as in Algorithm 1. When it is agent i 's turn, we determine i 's most preferred alternative a_i whose probability can still be increased (while also increasing the probability for alternatives a_j with $j < i$). Having determined a_i for every agent i , we find the maximal δ so that the probability for all a_i can be increased by at least δ . Note that the probability of every alternative a_i is increased by the same amount independently of how many agents nominate a_i .

This version of VER can address several classes of problems.

Probabilistic voting This model captures probabilistic voting (see, for example, Brandt, 2019) where voters have preferences over the alternatives A and an outcome is a probability distribution over the alternatives. That is, $X = \Delta(A) \subset [0, 1]^A$. In this context, VER coincides with the ESR rule of Aziz and Stursberg (2014).

Participatory budgeting The probabilistic voting setting can also be interpreted as determining the share of the budget allocated to each of the alternatives (Airiau et al., 2019). Since our model allows for arbitrary constraints, it can capture natural constraints such as enforcing lower bounds (reflecting the minimum funding required) on alternatives that get at least some funding.

Probabilistic allocation with bundles Let O be a set of objects and assume every agent i has a preference relation $\hat{\succsim}_i$ over *subsets* of O (see, for example, Chatterji and Liu, 2020). A deterministic allocation of objects to agents is an ordered N -partition of objects (which may include empty sets). Let A be the set of all such partitions. By $a(i)$ we denote the set of objects

agent i receives in the allocation a . The preference relation \succsim_i over A has $a \succsim_i b$ if and only if $a(i) \succsim_i b(i)$ for all $a, b \in A$. A probabilistic allocation is an element of the unit simplex $\Delta(A) \subset [0, 1]^A$ and $X \subset \Delta(A)$ specifies a set of feasible probabilistic allocations.

Two-sided probabilistic matching Let N_1 and N_2 be disjoint sets of agents. Each agent $i \in N_1$ has a preference relation $\hat{\succsim}_i$ over agents in N_2 and likewise for agents in N_2 . A deterministic matching is a subset μ of $N_1 \times N_2$ such that $(i, j), (i, j') \in \mu$ implies $j = j'$ and $(i, j), (i', j) \in \mu$ implies $i = i'$. Let A be the set of all deterministic matchings. The preference relation \succsim_i of $i \in N_1$ over A has $\mu \succsim_i \mu'$ if and only if $(i, j) \in \mu$ and $(i, j') \in \mu'$ with $j \hat{\succsim}_i j'$ or $(i, j) \in \mu'$ for no $j \in N_2$; preferences for agents in N_2 are defined analogously. A random matching is an element of the unit simplex $\Delta(A)$ and X specifies a set of feasible random matchings. We can apply VER to the problem with two-sided preferences and let both sides eat simultaneously. This approach is promising because many standard mechanisms for two-sided matching are asymmetric in that they treat the two sides differently. If we do not impose any stability constraints and both sides have dichotomous preferences (in which there are at most two indifference classes), then VER is equivalent to [Bogomolnaia and Moulin’s \(2004\)](#) egalitarian rule.

We assumed that agents have complete preference orders over alternatives. VER and its properties extend to preferences given by partial orders. The change that is required is the same that [Katta and Sethuraman \(2006\)](#) suggested for adapting the extended probabilistic serial rule to partial orders. Instead of trying to increase the probability for most preferred alternatives, agents try to increase the probability for those alternatives that are not strictly dominated by any other alternative.

Acknowledgements

The authors thank Fuhito Kojima, Barton Lee, Debasis Mishra, Hervé Moulin, Arunava Sen, and two anonymous referees for helpful comments. They also thank the participants of the following events where the paper was presented: Indian Statistical Institute Seminar Series; COMSOC Video Seminar; INFORMS Workshop on Market Design 2021; and the 1st IJCAI-PRICAI Workshop on Applied Mechanism Design. Florian Brandl acknowledges support by the Deutsche Forschungsgemeinschaft under grant BR 5969/1-1.

APPENDIX

A. Proofs From Section 4

Theorem 1. *For every non-empty and compact set of allocations X , Algorithm 1 terminates after at most $m \cdot n$ rounds with an allocation $p \in X$.*

Proof. First, we show that the optimization problem in Line 13 always has a solution. Second, we show that from each round to the next, either some agent is removed from the set of active

agents N' or some agent moves to a less preferred equivalence class.

Let k be the index of a round in the algorithm. For the first statement, we have to show that there exists (p, δ) that satisfies the constraints in Line 13 in round k . If $k = 0$, this is trivial since $\pi(i, E) = 0$ for all $i \in N$ and $E \in \mathcal{E}_i$ and X is non-empty. If $k > 0$, let (p, δ^*) be an optimal solution to the optimization problem computed in round $k - 1$. It follows from the definition of π^k (at the end of round $k - 1$) that $(p, 0)$ is a feasible point of the optimization problem in round k . Since X is compact, the problem has an optimal solution.

For the second statement, denote by \mathcal{F}_i^k the set of equivalence classes available to agent i in round k (cf. Line 8), by N'_k the set of agents so that $\mathcal{F}_i^k \neq \emptyset$, and by $E_i^k = \max_{\succsim_i} \mathcal{F}_i^k$ for $i \in N'$ their most preferred equivalence classes. Define \mathcal{F}_i^{k+1} , N'_{k+1} , and, for $i \in N'_{k+1}$, E_i^{k+1} similarly. Note that $N'_k \neq \emptyset$ since we would not have reached round $k + 1$ otherwise. Moreover, $N'_k \neq \emptyset$ implies that δ^* as computed in Line 13 in round k is strictly positive by definition of \mathcal{F}_i^k .

Now if $N'_{k+1} = \emptyset$, some agent is removed from the set of active agents and there is nothing left to show. So assume $N'_{k+1} \neq \emptyset$. We want to show that for $i \in N'$, $\mathcal{F}_i^{k+1} \subset \mathcal{F}_i^k$. If $E \in \mathcal{F}_i^{k+1}$, there is $p \in X$ such that $p(j, E') \geq \pi^{k+1}(j, E')$ for all $j \in N$ and $E' \in \mathcal{E}_j$, $p(i, E) > \pi^{k+1}(i, E)$, and $p(j, E_j) > \pi^{k+1}(j, E_j)$ for all $j \in N'_{k+1}$ with $j < i$. These properties of p , the definition of π^{k+1} , and the fact that $\delta^* > 0$ in round k imply that $E \in \mathcal{F}_i^k$. So we get $E_i^k \succsim_i E_i^{k+1}$ for all $i \in N'_{k+1}$. If $N'_{k+1} = N'_k$, it follows from the choice of δ^* that this preference is strict for at least one $i \in N'_{k+1}$. Otherwise $N'_{k+1} \subsetneq N'_k$. It follows that the algorithm terminates after at most $m \cdot n$ rounds. Clearly, the returned allocation p is in X . \square

Theorem 2. *When VER is applied to any non-empty and compact set of allocations X , it returns an allocation that is constrained LD-efficient among the allocations in X .*

Proof. Let \succsim_N be a preference profile and p be the allocation returned by VER for X and \succsim_N . For every allocation q and $t \geq 0$, let q^t be an allocation where every agent i receives a prefix (according to her preferences) of $q(i)$ summing to $\min\{t, q(i, O)\}$. Formally, $q^t(i, O) = \min\{t, q(i, O)\}$ and, for $E_i \in \mathcal{E}_i$, $q^t(i, E_i) > 0$ implies $q^t(i, E'_i) = q(i, E'_i)$ for all $E'_i \succ_i E_i$.

Assume there is an allocation $q \in X$ that lexicographically dominates p . Let

$$t^* = \max\{t \geq 0: p^t(i) \sim_i^{LD} q^t(i) \text{ for all } i \in N\}$$

Let i^* be the lexicographically first agent so that $q^{t^*+\epsilon}(i) \succ_i^{LD} p^{t^*+\epsilon}(i)$ for all $\epsilon > 0$.

Observe that $p^{t^*}(i^*, E_i) = p(i^*, E_i)$ whenever $p^{t^*}(i^*, E_i) > 0$. To see this, let E_i^* be i^* 's least preferred equivalence class those with $p^{t^*}(i^*, E_i) > 0$. It certainly holds that $p^{t^*}(i^*, E_i) = p(i^*, E_i)$ for all E_i with $E_i \succ_{i^*} E_i^*$ by definition of p^{t^*} . Now if $p(i^*, E_i^*) > p^{t^*}(i^*, E_i^*)$, let $0 < \epsilon < p(i^*, E_i^*) - p^{t^*}(i^*, E_i^*)$. We have $p^{t^*+\epsilon}(i^*, E_i) = q^{t^*+\epsilon}(i^*, E_i)$ for all $E_i \succ_{i^*} E_i^*$ by the choice of t^* . Moreover, by the choice of ϵ , $p^{t^*+\epsilon}(i^*, O) = q^{t^*+\epsilon}(i^*, O)$ and so $p^{t^*+\epsilon}(i^*, E_i^*) = q^{t^*+\epsilon}(i^*, E_i^*)$. But this contradicts $q^{t^*+\epsilon}(i^*) \succ_{i^*}^{LD} p^{t^*+\epsilon}(i^*)$. It follows that agent i^* moves to a less preferred equivalence class after securing t^* of equivalence classes at least as good as E_i^* . That is, there is a round k in Algorithm 1 such that $p^k(i, E_i) = \pi(i, E_i)$ for all $i \in N$ and $E_i \in \mathcal{E}_i$.

Let N' be the set of active agents at the end of round k , that is, all agents who increase their guarantee for some equivalence class in round k . For every $i \in N'$, let E_i be as determined by Line 8 in round k . Since p is the outcome of VER, it follows that $p(i, E_i) > \pi^k(i, E_i)$ for all $i \in N'$. Moreover, $q(i, E_i) > \pi^k(i, E_i)$ for all $i < i^*$ (by the choice of i^*) and, since $q^{t^*+\epsilon} \succ_{i^*}^{LD} p^{t^*+\epsilon}$, there is $E'_{i^*} \succ_i E_{i^*}$ such that $q(i^*, E'_{i^*}) > p(i^*, E'_{i^*}) \geq \pi^k(i^*, E'_{i^*})$. This contradicts the choice of E_{i^*} since q is a witness that i^* could have chosen E'_{i^*} instead of E_{i^*} . \square

Theorem 3. *If X is a non-empty finite union of polytopes, the outcome of VER can be computed in polynomial time in the number of inequalities used to describe X .*

Proof. Let $X = \bigcup_{r=1}^s P_r$ so that P_r is a polytope for each r . The only computationally non-trivial steps are determining the set \mathcal{F}_i in Line 8 and the computation of δ^* in Line 13.

First, observe that given a round k of the algorithm, an agent i , an equivalence class $E \in \mathcal{E}_i$, and a polytope P_r , the problem

$$\begin{aligned} \epsilon^* &= \max_{\epsilon, p} \epsilon \text{ s.t.} \\ p(j, E_j) &\geq \pi^k(j, E_j) + \epsilon && \forall j \in N' \text{ with } j < i \\ p(i, E) &\geq \pi^k(i, E) + \epsilon \\ p(j, E') &\geq \pi^k(j, E') && \forall j \in N \text{ and } E' \in \mathcal{E}_j \\ p &\in P_r \end{aligned}$$

can be solved in polynomial time in the number of linear inequalities used to describe P_r . If $\epsilon^* > 0$, agent i can increase her probability for E . By solving this problem for every r and every equivalence class $E \in \mathcal{E}_i$, we can determine the most-preferred equivalence class E_i of which agent i can increase her share in polynomial time in the number of inequalities used to describe X .

Second, the linear program

$$\begin{aligned} \delta_r^* &= \max_{\delta, p} \delta \text{ s.t.} \\ p(i, E_i) &\geq \pi^k(i, E_i) + \delta && \forall i \in N' \\ p(i, E) &\geq \pi^k(i, E) && \forall i \in N \text{ and } E \in \mathcal{E}_i \\ p &\in P_r \end{aligned}$$

can be solved in polynomial time for every P_r . Since $\delta^* = \max_r \delta_r^*$, the value of δ^* in Line 13 can be computed in polynomial time.

Theorem 1 ensures that the number of iterations of the while-loop in Line 4 is bounded by $m \cdot n$. \square

Theorem 4. *The VER allocation for a non-empty, compact, and convex set of allocations X is ordinally egalitarian among the allocations in X .*

Proof. Consider an allocation q that is ordinally egalitarian and an allocation p that is returned by VER. Consider $t^\uparrow(q) = (t_i^j(q))_{i,j}$ the signature vector of q and $t^\uparrow(p) = (t_i^j(p))_{i,j}$ the signature vector of p .

Consider that during the run of VER, lower bound constraints of the following form are added: $p(E_i^j) \geq \lambda$. The set of such constraints can alternatively be written as lower bounds on the upper contour set as follows: $p(\bigcup_{\ell=1}^j E_i^\ell) \geq \lambda'$. Equivalently, they can be written as $t_i^j \geq \lambda'$. At the start of round k , we denote by $g(i, k)$ the number of the first equivalence class $E_i^{g(i,k)}$ of agent i for which the lower bound has not been fixed.

We prove by induction on the rounds k of the algorithm that VER finds the largest s such that $t_i^{g(i,k)} \geq s$ and that the following values are present in $t^\uparrow(q)$: $\min_{i' \in N} p(\bigcup_{\ell=1}^{g(i',k)} E_{i'}^\ell)$ as well as $p(\bigcup_{\ell=1}^{k'} E_{i'}^\ell)$ for $i \in N$ and $k' < g(i', k)$.

For $k = 0$, in the first round, VER tries to maximize the lower bound on $E_i^{g(i,1)}$ for all $i \in N$. Suppose that $p^1(E_i^{g(i,1)}) \geq \lambda'$ for all $i \in N$. Then λ' is by definition δ^* as computed by VER. It also follows that for all $i \in N$, $p^1(E_i^j) = 0$ for all $j < g(i, 1)$. Hence, we have established that the minimum non-zero entry in the vector $t^\uparrow(p)$ is the same as the minimum non-zero entry in vector $t^\uparrow(q)$. By convexity of the feasible region it follows that it is not possible to have an allocation in which $p(\bigcup_{\ell=1}^j E_i^\ell) > 0$ for any $j < g(i, 1)$. Therefore, the corresponding entries $t_i^\ell = 0$ are also entries in $t^\uparrow(q)$.

Now suppose k rounds have passed. By the induction hypothesis, for each $k' \leq k$, $t_i^{g(i,k')}$ is present in the vector $t^\uparrow(q)$. We also note that VER has fixed a weight for each of $\bigcup_{\ell=1}^j E_i^\ell$ where $j < g(i, k+1)$. At this point, VER computes the largest δ that can be additionally guaranteed for each $E_i^{g(i,k+1)}$. Equivalently, it computes the largest s such that $p(\bigcup_{\ell=1}^{g(i,k+1)} E_i^\ell) \geq s$ for all $i \in N$. Then $\min_{i' \in N} p(\bigcup_{\ell=1}^{g(i',k+1)} E_{i'}^\ell)$ is the optimized weight of the next heavy upper contour set. Hence $\min_{i' \in N} p(\bigcup_{\ell=1}^{g(i',k+1)} E_{i'}^\ell)$ is present in $t^\uparrow(q)$. \square

Theorem 5. *Let X be a non-empty and compact set of allocations and \succsim_N be a preference profile. If X is symmetric and convex for all $i, i' \in N$ with $\succsim_i = \succsim_{i'}$, then the outcome of VER satisfies equal treatment of equals.*

Proof. Let p be the outcome of VER for the set X and the profile \succsim_N . We prove by induction that for each round k of Algorithm 1, $\pi^k(i, E) = \pi^k(i', E)$ for all $E \in \mathcal{E}_i = \mathcal{E}_{i'}$. It then follows that $p(i, E) = p(i', E)$ for all $E \in \mathcal{E}_i$.

For $k = 0$, this is clear since $\pi^0(j, E) = 0$ for all $j \in N$ and $E \in \mathcal{E}_j$. Now let k be arbitrary and assume that $\pi^k(i, E) = \pi^k(i', E)$ for all $E \in \mathcal{E}_i$. Let N' be the set active agents at the beginning of round k . We show that $\mathcal{F}_i = \mathcal{F}_{i'}$, where the \mathcal{F}_j are determined as in Line 8. Assume that $i < i'$. Since X is symmetric for i and i' , it follows that $\mathcal{F}_{i'} \subseteq \mathcal{F}_i$. If $\mathcal{F}_i = \emptyset$, we are done. Otherwise, let $j = \max\{j' \in N' : j' < i' \text{ and } \mathcal{F}_{j'} \neq \emptyset\}$ (note that $j \geq i$). By definition of E_j , there is an allocation $p \in X$ such that $p(j', E) \geq \pi^k(j', E)$ for all $j' \in N$ and $E \in \mathcal{E}_{j'}$ and $p(j', E_{j'}) > \pi^k(j', E_{j'})$ for all $j' \in N'$ with $j' \leq j$ and $\mathcal{F}_{j'} \neq \emptyset$. Since X is symmetric and convex for i and i' , the allocation $p' = \frac{1}{2}p + \frac{1}{2}p^{(ii')}$ is in X . Thus defined, p' satisfies the constraints on p above. For $j' \neq i, i'$, this is obvious since $p'(j') = p(j')$. For i and i' , it follows from

$p(i, E_i) > \pi^k(i, E_i)$ and $\pi^k(i, E) = \pi^k(i', E)$. Clearly, $p'(i) = p'(i')$ and so $p'(i', E_i) > \pi^k(i', E_i)$. It follows that $E_i \in \mathcal{F}_{i'}$. Thus, $E_i = E_{i'}$ and so $\pi^{k+1}(i, E) = \pi^{k+1}(i', E)$ for all $E \in \mathcal{E}_i$. \square

Theorem 6. *Let X be a non-empty, compact, and convex set of allocations and $p \in X$. Then p is constrained LD-efficient with respect to X if and only if all agents are indifferent between p and the outcome of some instance of one-at-a-time-VER on the set X .*

Proof. Our arguments in the proof of Theorem 2 go through to show that the outcome of any instance of one-at-a-time-VER is constrained LD-efficient.

For the converse, assume that p is constrained LD-efficient with respect to X . Our arguments are similar to those used in the proof of Theorem 8 by Aziz and Stursberg (2014). We show that if π is an allocation to equivalence classes so that $p(i, E_i) \geq \pi(i, E_i)$ for all $i \in N$ and $E_i \in \mathcal{E}_i$ with at least one strict inequality, then there is an agent i^* so that the most-preferred equivalence class E_{i^*} for which i^* can increase her probability has $p(i^*, E_{i^*}) > \pi(i^*, E_{i^*})$. Then, by allowing i^* to eat, we can get closer to our target allocation p and the claim follows.

Suppose we have π as above. For every $i \in N$, let $\mathcal{F}_i = \{E \in \mathcal{E}_i : \text{there is } q \in X \text{ so that } q(j, E_j) \geq \pi(j, E_j) \text{ for all } j \in N \text{ and } E_j \in \mathcal{E}_j \text{ and } q(i, E) > \pi(i, E)\}$. Let $N' = \{i \in N : \mathcal{F}_i \neq \emptyset\}$ and for $i \in N'$, let $E_i = \max_{\succsim_i} \mathcal{F}_i$. Note that N' is non-empty since we assume that $p(i, E) > \pi(i, E)$ for some $i \in N$ and $E \in \mathcal{E}_i$. Since X is convex, we can find $p^* \in X$ so that $p^*(i, E) \geq \pi(i, E)$ for all $i \in N$ and $E \in \mathcal{E}_i$ and $p^*(i, E_i) > \pi(i, E_i)$ for all $i \in N'$.

Now if $p(i, E'_i) > \pi(i, E'_i)$ for $E'_i \in \mathcal{E}_i$, then $E_i \succsim_i E'_i$ since p is a witness that $E'_i \in \mathcal{F}_i$. Since p is LD-efficient, it is not LD-dominated by p^* . So either $p(i) \sim_i^{LD} p^*(i)$ for all $i \in N$ or there is $i \in N$ such that $p(i) \succ_i^{LD} p^*(i)$. In the first case we can choose $i^* \in N'$ arbitrarily. In the second case, choose $i \in N$ so that $p(i) \succ_i^{LD} p^*(i)$. So we can find E'_i with $p(i, E'_i) > p^*(i, E'_i)$ and $p(i, E''_i) = p^*(i, E''_i)$ for all $E''_i \succ_i E'_i$. It follows that $E'_i \in \mathcal{F}_i$ and so $i \in N'$ and $E_i \succsim_i E'_i$. The latter implies that $p(i, E_i) \geq p^*(i, E_i) > \pi(i, E_i)$. We can thus choose $i^* = i$. \square

B. Proofs From Section 5

Corollary 1. *If S is ex-ante or ex-post stability, S -VER returns an S -stable allocation that is SD-efficient among S -stable allocations and satisfies limited equal treatment of equals.*

Proof. Stability and constrained SD-efficiency of the VER allocation follow from Theorems 1 and 2.

To show that VER satisfies limited equal treatment of equals, let $i, i' \in N$ such that $\succsim_i = \succsim_{i'}$, $\text{cap}(i) = \text{cap}(i')$, and $i \sim_o i'$ for all $o \in O$. We prove that the set of S -stable allocations is symmetric and convex for $\{i, i'\}$. Then Theorem 5 yields the desired conclusion. Symmetry is obvious for both stability notions. Also, for ex-post stability, the set of stable allocations is convex and thus convex for $\{i, i'\}$. The set of ex-ante stable allocations is not in general convex. It is convex for $\{i, i'\}$, however, as we show now.

Let p be ex-ante stable and $\lambda \in [0, 1]$; let $q = \lambda p + (1 - \lambda)p^{(ii')}$. Note that $q(j) = p(j)$ for all $j \neq i, i'$. Thus, it suffices to check envy-freeness for pairs of agents in $\{i, i'\} \times N \setminus \{i, i'\}$. Let

$j \in N \setminus \{i, i'\}$. First, if $i \succ_o j$ and $q(j, o) > 0$, then $p(j, o) > 0$ since $p(j) = p^{(i')}(j)$. Thus, $\sum_{o' \succ_{i'o} o} p(i, o') = \sum_{o' \succ_{i'o} o} p(i', o') = \text{cap}(i) = \text{cap}(i')$. It follows that $\sum_{o' \succ_{i'o} o} q(i, o') = \text{cap}(i)$ and similarly for i' . Second, if $j \succ_o i$ and $q(i, o) > 0$ (or, equivalently, $q(i', o) > 0$), then without loss of generality, $p(i, o) > 0$. Hence, $\sum_{o' \succ_{jo} o} p(j, o') = \sum_{o' \succ_{jo} o} p^{(i')}(j, o') = \text{cap}(j)$. We get $\sum_{o' \succ_{jo} o} q(j, o') = \text{cap}(j)$ as desired. Together with the fact that $q(j) = p(j)$ for $j \neq i, i'$, this shows that q is ex-ante stable. \square

Corollary 2. *Assume that preferences and priorities are strict. Then, if S is ex-ante stability or ex-post stability, S -VER returns the agent optimal deterministic stable allocation.*

Proof. We use that the sets of ex-ante stable allocations and of ex-post stable allocations are lattices when join and meet are defined via stochastic dominance. This has been shown by Alkan and Gale (2003, Theorem 5) for ex-ante stability (and a more general class of preferences) and by Juárez et al. (2022) for ex-post stability. The upper bound (with respect to the join operation) of both lattices is the agent-optimal deterministic stable allocation, which thus stochastically dominates every other ex-post or ex-ante stable allocation according the agents' preferences. Since by Corollary 1 the outcome of S -VER is SD -efficient among S -stable allocations, it follows that S -VER returns the agent-optimal deterministic stable allocation when S is ex-ante stability or ex-post stability. \square

Proposition 1. *Assume that the agents' preferences are strict. If S is ex-ante stability, S -VER returns an allocation that is two-sided SD -efficient.*

Proof. Let (\succ_N, \succ_O) be a profile and p be an allocation that is ex-ante stable and SD -efficient among ex-ante stable allocations. By Corollary 1, it suffices to show that p is two-sided SD -efficient.

If p is not two-sided SD -efficient, there exists a cycle $(i_0, o_0), (i_1, o_1), \dots, (i_k, o_k)$ such that $(i_0, o_0) = (i_k, o_k)$ and for all $l = 1, \dots, k$, $o_l \succ_{i_l} o_{l-1}$, $i_{l-1} \succ_{o_{l-1}} i_l$, and $p(i_l, o_{l-1}) > 0$ (cf. Dogan and Yildiz, 2016, Proposition 1). For $\epsilon \in (0, \min\{p(i, o) : i \in N \text{ and } o \in O\})$, let q be equal to p except that $q(i_l, o_l) = p(i_l, o_l) + \epsilon$ and $q(i_l, o_{l-1}) = p(i_l, o_{l-1}) - \epsilon$ for all $l = 1, \dots, k$. Thus, agent i_l passes a fraction of o_{l-1} on to agent i_{l-1} . Note that q stochastically dominates p for both sides and at least one agent $i \in N$ strictly prefers q to p since preferences are strict.

We show that q is ex-ante stable. Let $i, j \in N$ and $o \in O$ with $i \succ_o j$ and $q(j, o) > 0$. If $p(j, o) > 0$, ex-ante stability of p implies $p(i, o') = 0$ for all o' with $o \succ_i o'$. Since $q(i) \succ_i^{SD} p(i)$, it follows that $q(i, o') = 0$ for all o' with $o \succ_i o'$. Otherwise, $j = i_{l-1}$ and $o = o_{l-1}$ for some $l = 1, \dots, k$. Now $p(i_l, o_{l-1}) \geq \epsilon > 0$ and $i_{l-1} \succ_{o_{l-1}} i_l$. It follows that $i \succ_{o_{l-1}} i_l$. Since p is ex-ante stable, $p(i, o') = 0$ for all o' with $o_{l-1} \succ o'$. And again, since $q(i) \succ_i^{SD} p(i)$, it follows that $q(i, o') = 0$ for all o' with $o_{l-1} \succ o'$. So q is ex-ante stable. Thus, p is not SD -efficient among ex-ante stable allocations, which contradicts the assumption. \square

Proposition 2. *There may be no allocation that is ex-post stable, weakly SD -envy-free, and SD -efficient among ex-post stable allocations (if there are at least three agents and three objects).*

Proof. Assume all capacities are equal to 1 and consider the following profile.

$$\begin{array}{lcl} \succ_1: & a & b & c & \succ_a: & 3 & 1, 2 \\ \succ_2: & a & c & b & \succ_b: & 1, 2 & 3 \\ \succ_3: & b & a & c & \succ_c: & 1, 2 & 3 \end{array}$$

Then, the only deterministic stable allocations are $p = \{(1, a), (2, c), (3, b)\}$ and $q = \{(1, b), (2, c), (3, a)\}$. Note that p Pareto dominates q . Hence, no convex combination of p and q with positive weight on q is ex-post efficient (and, thus, SD -efficient) among ex-post stable allocations. The unique ex-post stable allocation that is SD -efficient among ex-post stable allocations is thus p . Since p is not SD -envy-free, the result follows. \square

Proposition 3. *If S is ex-ante stability or ex-post stability, no mechanism satisfies weak SD -strategyproofness, S -stability, and S -constrained efficiency even if preferences are strict (and there are at least three agents and three objects).*

Proof. For all profiles in the proof, the set of ex-ante stable allocations and the set of ex-post stable allocations will coincide. Thus, it will prove the statement for both stability notions at once. Moreover, we assume all capacities are equal to 1 throughout.

Let f be a mechanism that is weakly SD -strategyproof, S -stable, and S -constrained efficient. Consider the following profile P .

$$\begin{array}{lcl} \succ_1: & b & c & a & \succ_a: & \{1, 2\} & 3 \\ \succ_2: & b & c & a & \succ_b: & 3 & \{1, 2\} \\ \succ_3: & a & b & c & \succ_c: & 3 & \{1, 2\} \end{array}$$

There are two S -stable and S -constrained efficient deterministic allocations.

$$\begin{array}{ccc} & a & b & c & & a & b & c \\ 1 & \left(\begin{array}{ccc} 0 & 1 & 0 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{array} \right) & & & 1 & \left(\begin{array}{ccc} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \end{array} \right) \\ 2 & & & & 2 & & & \\ 3 & & & & 3 & & & \end{array}$$

Thus, all S -stable and S -constrained efficient allocations are of the following form. (Here we use that the set of S -stable allocations in the same for both stability notions.)

$$p^\alpha = \begin{array}{ccc} & a & b & c \\ 1 & \left(\begin{array}{ccc} 0 & 1 - \alpha & \alpha \\ 0 & \alpha & 1 - \alpha \\ 1 & 0 & 0 \end{array} \right) \\ 2 & & & \\ 3 & & & \end{array}$$

for some $\alpha \in [0, 1]$. Suppose $f(P) = p^\alpha$ for some $\alpha > 0$. Then, agent 1 can misreport by swapping c and a resulting in the preferences \succ'_N below.

$$\begin{array}{lcl} \succ'_1: & b & a & c \\ \succ'_2: & b & c & a \\ \succ'_3: & a & b & c \end{array}$$

S -stability and S -constrained efficiency imply that $f((\succ'_N, \succ_O)) = p^0$.

Note that $p^0(1) \succ_i^{SD} p^\alpha$ and so agent 1 can successfully manipulate in the profile P , which contradicts strategyproofness of f . If $1 - \alpha > 0$, we can use a symmetric argument where agent 2 misreports. \square

C. Comparison of Probabilistic Allocation Mechanisms Under Priorities

We discuss other mechanisms presented in the literature. Since most of them have only been defined for agents with strict preferences, we assume that preferences are strict for the comparison. We have already discussed the probabilistic serial rule in the introduction. Except for the first mechanism (random priority), all the other mechanism are extensions of the deferred acceptance algorithm, which is typically defined for strict preferences and strict priorities (Gale and Shapley, 1962; Roth, 2008) and returns a stable allocation.

Random Priority (or random serial dictatorship) chooses an ordering of the agents uniformly at random and then lets each agent pick her most preferred object among the ones remaining in that order (Bogomolnaia and Moulin, 2001; Aziz et al., 2013a). The rule does not take into account the priorities of the objects. For the basic assignment problem, random priority is known to be strategyproof. It satisfies equal treatment of equals but is not SD -efficient or SD -envy-free (Bogomolnaia and Moulin, 2001). If the priority order is chosen deterministically, then SD -efficiency is regained, but equal treatment of equals is lost.

Deferred acceptance with lexicographic tie-breaking A simple adaptation of the deferred acceptance algorithm to the case of weak priorities is to break the ties and then run deferred acceptance. If the tie-breaking is pre-determined, for example lexicographically over agents, and thus does not depend on the agents' preferences, the resulting mechanism is strategyproof. Moreover, it returns a deterministic allocation that is stable and hence ex-ante stable. However, it is not SD -efficient even among deterministic stable allocations. Like any other mechanism that returns deterministic allocations, it also violates limited equal treatment of equals.

Deferred acceptance with random tie-breaking Another natural approach is to break ties in the priorities uniformly at random and then run deferred acceptance. Afacan (2018) referred to this extension as *probabilistic deferred acceptance*. To compute the probabilistic allocation, we need to take the mean of the outcomes over all possible tie-breakings. Under flat priorities, probabilistic deferred acceptance is equivalent to random priority. Since the latter is well-known to violate SD -efficiency, it follows that probabilistic deferred acceptance does not satisfy constrained SD -efficiency.

| | S -VER | probabilistic serial | random priority | deferred acceptance (lexicographic) | deferred acceptance (uniform) | fractional deferred acceptance and trading |
|-----------------------------------|----------|-------------------------|--------------------|--|----------------------------------|---|
| S -stability | + | - | - | + ¹³ | + | + |
| S -constrained SD -efficiency | + | + | - | - | - | - ¹⁴ |
| SD -efficiency | - | + | - | - | - | - |
| SD -envy freeness | - | + | - | - | - | - |
| weak SD -envy freeness | - | + | + | - | - | - |
| limited equal treatment of equals | + | + | + | + | - | + |

Table 2: Summary of the properties satisfied by the allocation mechanisms discussed in Appendix C. Here, S is either ex-ante stability or ex-post stability. To make the rules comparable, we assume that preferences are strict.

Fractional deferred acceptance and trading Kesten and Unver (2015) presented the fractional deferred acceptance and trading mechanism. It returns an allocation that is strongly ex-ante stable and SD -efficient constrained to the set of strongly ex-ante stable allocations. It is different from the probabilistic serial rule under flat priorities and thus also from VER.

Constrained probabilistic serial Afacan (2018) introduced the *constrained probabilistic serial* rule, which is an adaptation of the probabilistic serial rule that obtains a claimwise stable allocation (see Footnote 12). Because of the simplicity of the constraints imposed by claimwise stability, this algorithm does not need to look ahead to check when an agent should stop eating an object. Suppose an agent i starts eating object o . At that point, we put an upper bound on all agents j who have a lower priority for o than i . If j had been eating o all the time while i was eating more preferred objects, we can stop j from eating more. If j had been eating o only part of the time, we put a limit on how much j can eat o . The upper limit is equal to the amount of time i was eating objects weakly more preferred objects. The constrained probabilistic serial rule can be viewed as a careful version of the probabilistic serial rule that handles stability constraints dynamically. By contrast, VER for claimwise stability makes look-ahead checks to see how much of an object an agent can eat before a stability violation becomes unavoidable. One can show that in this case, the look-ahead checks are not necessary and both mechanisms coincide.

Table 2 summarizes the properties satisfied by various mechanisms.

¹³The deferred acceptance algorithm with lexicographic tie-breaking satisfies ex-post stability. It violates ex-ante stability, however.

¹⁴The fractional deferred acceptance and trading algorithm violates S -constrained SD -efficiency when S is ex-post stability. It satisfies SD -efficiency constrained to the set of ex-ante stable allocations, however.

References

- A. Abdulkadiroğlu and T. Sönmez. School choice: A mechanism design approach. *American Economic Review*, 93(3):729–747, 2003.
- M. O. Afacan. The object allocation problem with random priorities. *Games and Economic Behavior*, 110:71–89, 2018.
- S. Airiau, H. Aziz, I. Caragiannis, J. Kruger, J. Lang, and D. Peters. Portioning using ordinal preferences: Fairness and efficiency. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*. IJCAI, 2019.
- M. Akbarpour and A. Nikzad. Approximate random allocation mechanisms. *The Review of Economic Studies*, 2020. Forthcoming.
- A. Alkan and D. Gale. Stable schedule matching under revealed preference. *Journal of Economic Theory*, 112(2):289–306, 2003.
- I. Ashlagi, A. Nikzad, and A. Romm. Assigning more students to their top choices: A comparison of tie-breaking rules. *Games and Economic Behavior*, 115:167–187, 2019a.
- I. Ashlagi, A. Saberi, and A. Shameli. Assignment mechanisms under distributional constraints. *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 229–240, 2019b.
- I. Ashlagi, A. Saberi, and A. Shameli. Assignment mechanisms under distributional constraints. *Operations Research*, 68:467–479, 2020.
- S. Athanassoglou and J. Sethuraman. House allocation with fractional endowments. *International Journal of Game Theory*, 40(3):481–513, 2011.
- H. Aziz. A probabilistic approach to voting, allocation, matching, and coalition formation. In J.-F. Laslier, H. Moulin, R. Sanver, and W. S. Zwicker, editors, *The Future of Economic Design*. Springer-Verlag, 2019.
- H. Aziz and B. Klaus. Random matching under priorities: Stability and no envy concepts. *Social Choice and Welfare*, 53(2):213–259, 2019.
- H. Aziz and P. Stursberg. A generalization of probabilistic serial to randomized social choice. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence (AAAI)*, pages 559–565. AAAI Press, 2014.
- H. Aziz, F. Brandt, and M. Brill. The computational complexity of random serial dictatorship. *Economics Letters*, 121(3):341–345, 2013a.
- H. Aziz, F. Brandt, and P. Harrenstein. Pareto optimality in coalition formation. *Games and Economic Behavior*, 82:562–581, 2013b.

- I. Balbuzanov. Constrained random matching. 2019. Working paper.
- P. Biró, T. Fleiner, R. W. Irving, and D. F. Manlove. The college admissions problem with lower and common quotas. *Theoretical Computer Science*, 411(34–36):3136–3153, 2010.
- A. Bogomolnaia. Random assignment: Redefining the serial rule. *Journal of Economic Theory*, 158:308–318, 2015.
- A. Bogomolnaia. The most ordinally egalitarian of random voting rules. *Journal of Public Economic Theory*, 20(2), 2017.
- A. Bogomolnaia and M. O. Jackson. The stability of hedonic coalition structures. *Games and Economic Behavior*, 38(2):201–230, 2002.
- A. Bogomolnaia and H. Moulin. A new solution to the random assignment problem. *Journal of Economic Theory*, 100(2):295–328, 2001.
- A. Bogomolnaia and H. Moulin. Random matching under dichotomous preferences. *Econometrica*, 72(1):257–279, 2004.
- A. Bogomolnaia, H. Moulin, and R. Stong. Collective choice under dichotomous preferences. *Journal of Economic Theory*, 122(2):165–184, 2005.
- F. Brandt. Rolling the dice: Recent results in probabilistic social choice. In U. Endriss, editor, *Trends in Computational Social Choice*, chapter 1, pages 3–26. AI Access, 2017.
- F. Brandt. Collective choice lotteries: Dealing with randomization in economic design. In J.-F. Laslier, H. Moulin, R. Sanver, and W. S. Zwicker, editors, *The Future of Economic Design*, Studies in Economic Design, pages 51–56. Springer-Verlag, 2019.
- E. Budish, Y.-K. Che, F. Kojima, and P. Milgrom. Designing random allocation mechanisms: Theory and applications. *American Economic Review*, 103(2):585–623, 2013.
- I. Caragiannis, A. Filos-Ratsikas, P. Kanellopoulos, and R. Vaish. Stable fractional matchings. In *Proceedings of the 20th ACM Conference on Economics and Computation (ACM-EC)*, pages 21–39, 2019.
- S. Chatterji and P. Liu. Random assignments of bundles. *Journal of Mathematical Economics*, 87:15–30, 2020.
- W. J. Cho and B. Dogan. Equivalence of efficiency notions for ordinal assignment problems. *Economics Letters*, 146:8–12, 2016.
- D. Delacrétaz. Processing reserves simultaneously. September 2020.
- B. Dogan and K. Yildiz. Efficiency and stability of probabilistic assignments in marriage problems. *Games and Economic Behavior*, 95:47–58, 2016.

- F. Echenique, A. Miralles, and J. Zhang. Constrained pseudo-market equilibrium. Papers 1909.05986, arXiv.org, Sept. 2019. URL <https://ideas.repec.org/p/arx/papers/1909.05986.html>.
- A. Erdil and H. Ergin. What’s the matter with tie-breaking? Improving efficiency in school choice. *American Economic Review*, 98(3):669–89, 2008.
- A. Erdil and H. Ergin. Two-sided matching with indifferences. *Journal of Economic Theory*, 171:268–292, 2017.
- S. Fujishige, Y. Sano, and P. Zhan. The random assignment problem with submodular constraints on goods. In *Proceedings of the 19th ACM Conference on Economics and Computation (ACM-EC)*, volume 6, pages 3:1–3:28, 2018.
- D. Gale and L. S. Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.
- X. Han. *Resource Allocation with Priorities*. PhD thesis, 2017. URL <https://search.proquest.com/docview/1917417439?accountid=12763>.
- Y. He, A. Miralles, M. Pycia, and J. Yan. A pseudo-market approach to allocation with priorities. *American Economic Journal: Microeconomics*, 10(3):272–314, 2018.
- A. Hylland and R. Zeckhauser. The efficient allocation of individuals to positions. *The Journal of Political Economy*, 87(2):293–314, 1979.
- M. O. Jackson and A. Wolinsky. A strategic model of social and economic networks. *Journal of Economic Theory*, 71(1):44–74, 1996.
- N. Juárez, P. Neme, and J. Oviedo. Lattice structure of the random stable set in many-to-many matching markets. *Games and Economic Behavior*, 132:255–273, 2022.
- Y. Kamada and F. Kojima. Efficient matching under distributional constraints: Theory and applications. *The American Economic Review*, 105(1):67–99, 2015.
- Y. Kamada and F. Kojima. Recent developments in matching with constraints. *The American Economic Review*, 107(5):200–204, 2017a.
- Y. Kamada and F. Kojima. Stability concepts in matching under distributional constraints. *Journal of Economic Theory*, 168:107–142, 2017b.
- A.-K. Katta and J. Sethuraman. A solution to the random assignment problem on the full preference domain. *Journal of Economic Theory*, 131(1):231–250, 2006.
- O. Kesten. School choice with consent. *The Quarterly Journal of Economics*, 125(3):1297–1348, 2010.

- O. Kesten and U. Ünver. A theory of school choice lotteries. *Theoretical Economics*, 10:543–595, 2015.
- F. Kojima. Random assignment of multiple indivisible objects. *Mathematical Social Sciences*, 57(1):134–142, 2009.
- F. Kojima, A. Tamura, and M. Yokoo. Designing matching mechanisms under constraints: An approach from discrete convex analysis. *Journal of Economic Theory*, 176:803–833, 2018.
- P. A. Pathak, T. Sönmez, M. U. Ünver, and M. B. Yenmez. Fair Allocation of Vaccines, Ventilators and Antiviral Treatments: Leaving No Ethical Value Behind in Health Care Rationing. Boston College Working Papers in Economics 1015, Boston College Department of Economics, July 2020. URL <https://ideas.repec.org/p/boc/bocoec/1015.html>.
- A. E. Roth. The economics of matching: Stability and incentives. *Mathematics of Operations Research*, 7(4):617–628, 1982.
- A. E. Roth. Deferred acceptance algorithms: History, theory, practice, and open questions. *International Journal of Game Theory*, 36:537–569, 2008.
- A. E. Roth and M. A. O. Sotomayor. *Two-Sided Matching: A Study in Game Theoretic Modelling and Analysis*. Cambridge University Press, 1990.
- A. E. Roth, U. G. Rothblum, and J. H. Vande Vate. Stable matchings, optimal assignments, and linear programming. *Mathematics of Operations Research*, 18(4):803–828, 1993.
- T. Sönmez. Strategy-proofness and essentially single-valued cores. *Econometrica*, 67(3):677–689, 1999.
- L.-G. Svensson. Queue allocation of indivisible goods. *Social Choice and Welfare*, 11:323–330, 1994.
- L.-G. Svensson. Strategy-proof allocation of indivisible goods. *Social Choice and Welfare*, 16(4):557–567, 1999.
- C.-P. Teo and J. Sethuraman. The geometry of fractional stable matchings and its applications. *Mathematics of Operations Research*, 23(4):874–891, 1998.
- O. Yilmaz. The probabilistic serial mechanism with private endowments. *Games and Economic Behavior*, 69(2):475–491, 2010.