

Fair Congested Assignment Problem*

Anna Bogomolnaia[†] and Hervé Moulin[‡]

February 2024

Abstract

We must assign n agents to m posts subject to negative congestion: what assignment is fair and efficient?

If congestion is anonymous (each agent adds one unit) it is always possible to assign each agent to one of their top n out of the $n \times m$ feasible allocations. This ordinal interpretation of Ex Ante fairness can be adjusted if congestion is weighted (agent-specific).

An assignment is *Competitive* if I don't want to move to an empty post, or to an occupied one *at its current congestion level*. If it exists the competitive assignment is essentially unique, efficient and Ex Ante fair.

Among agents endowed with cardinal vNM utilities we can randomise the selection of our assignment. Under anonymous congestion every problem has a unique competitive congestion profile, implemented by a mixture of deterministic assignments rounding up or down the competitive congestion, and approximately Ex Ante fair, efficient and welfare equivalent. Some of these properties are lost under weighted congestion.

Keywords: assignment, congestion, guarantee, competitive, envy free, free mobility game

JEL classification: D47, D63

* *Acknowledgments:* We thank Felix Fischer, Ioannis Caragiannis, Igal Milchtaich and Sinan Ertemel for critical comments and references.

[†]Adam Smith Business School, University of Glasgow and Centre d'Economie de la Sorbonne, CNRS, Paris

[‡]Adam Smith Business School, University of Glasgow

1 introduction

Congestion affects the allocation of any commodity consumed under partial rivalry: traffic on roads and the internet, public utilities and private services, and so on: for just over a century it is a staple of microeconomic analysis. Pigou's work on the taxation of externalities([31]) inspired the search for the welfare-optimal taxation of congestion, for instance peak load pricing of utilities ([32], [8]), Vickrey's optimal toll road ([39]) and the far reaching Vickrey-Clarke-Groves pricing mechanisms ([18]).

A different and equally outstanding operations research and game theory literature starts with Wardrop's transportation model ([14]): congestion is analysed in the decentralised *free mobility* regime: each agent chooses her service (what route to use, which store to visit, etc..) and the congestion results from these non cooperative interactions (typically without any transfer of money). Such games include the seminal instance of potential games ([33], [26]) and the first definition of the price of anarchy¹ ([24], [34]).

Our model uses the same physical description of assignments as a familiar congestion game (e. g. [25]): there is a set of heterogenous items subject to congestion that we call "posts"; each agent must be assigned to a post, and each post can take any amount of congestion, from remaining empty to hosting everyone. The only restriction on agents' preferences is that at any given post they decrease strictly in the amount of congestion. Critically cash transfers (taxes) are ruled out: agents cannot be compensated to accept an inferior post, or pay to be assigned at a popular one. Our viewpoint is purely normative, non strategic: which of the efficient assignments should we call fair?

Formal research on this aspect of congestion is scarce and quite recent (see details in subsection 1.2). Yet moneyless assignment problems where congestion is an important consideration are easy to find: the allocation of jobs to busy heterogeneous servers, of workers to shared office spaces, patients to hospitals, students to crowded classes, students to crowded schools, or messages to routes in a centralised communication network. In those examples the decentralised free mobility "choose your own post" approach is clearly impractical (and possibly inefficient). It is also unpalatable because it gives an unfair advantage to agents who happen to play the game better: see the discussion of the perverse strategic implications of the Boston

¹Measuring the welfare loss in the worst Nash equilibrium of the game; see e. g. [17], [35].

school assignment mechanism ([15]).

Matching and assignment models involving “many agents to one position” routinely incorporate *hard* congestion constraints in the form of upper or lower bounds on the filling of each post: maximal capacity of a class ([9]) or a school ([1]), minimal or maximal quotas for some subsets of students ([3], [19], [27], [21]) etc.. But when these constraints do not bite or are altogether absent, *soft* congestion still impacts agents’ welfare and choices. Parents will accept more crowded classes if the school’s academic context is better, or vice versa: the recent paper [30] on congested school choice (on which more in the next section) offers clear evidence of this point.

Our definition of fairness, in line with seven decades of microeconomic literature on the division of private commodities (e. g. [28]), is two-fold. First, Ex Ante fairness focuses on the worst case welfare level that can be guaranteed to each agent, based on their own trade-offs between posts and congestion but no other information about the preferences of the other potentially adversarial participants. Second, Ex Post fairness adapts the familiar Envy-Freeness property ([16], [38]) under the name of *Competitiveness*, to account for the fact that the congestion itself plays the role of a price.

Our results. We identify the powerful Ex Ante test of “top-fairness” relying only on ordinal preferences; it is always feasible and often met by only a handful of assignments. In the deterministic model the Competitiveness property is not always feasible, but if it is it selects a unique, fair and efficient solution. When we randomise assignments and the agents have cardinal vNM utilities, there is always a unique *fractional* competitive congestion profile, implemented by approximately fair and efficient deterministic assignments.

1.1 overview of the paper

The deterministic model with ordinal preferences, relevant until section 5, is defined in the short section 2. If congestion is anonymous (unweighted) an agent’s allocation is a pair (a, s_a) where a is the post and s_a the number of agents at a . If each agent i brings w_i units of congestion to any post an allocation is (a, w_{S_a}) where w_{S_a} is the total weight of the set S_a of the agents assigned to a . In both cases preferences decrease strictly in the congestion at a .

In section 3 we have m posts, n agents and congestion is anonymous; this gives

$m \times n$ distinct allocations. An assignment is top- n -fair if every agent i 's allocation is one of i 's top n allocations (Definition 1). Top- n -fairness is always feasible and compatible with efficiency (Pareto optimality); it cannot be improved: Proposition 1.

In section 4 congestion is weighted and the definition of top-fairness is more subtle. At the Ex Ante stage agent i knows only their own weight w_i and the total weight W of all agents, but not their number: so i 's worst case allows for any finite split of $W - w_i$. The set of i 's feasible allocations is the union of m intervals $[w_i, W]$, of size $m \times (W - w_i)$. An assignment is top- $\frac{1}{m}$ -fair if each agent i gets one of their top $(W - w_i)$ allocations (Definition 2).² Such assignments always exist (Proposition 2) and this guarantee cannot be improved.

Competitiveness in the ordinal model is the subject of section 5. Under anonymous congestion (subsection 5.1) an assignment is competitive if each agent i weakly prefers their allocation (a, s_a) to any other allocation (x, s_x) at this assignment, including to $(y, 1)$ if the post y is empty ($s_y = 0$) (Definition 3). Taking the largest of s_x and 1 as the congestion price of post x , the competitive demand for post a contains s_a (equals s_a if each demand is single-valued). If it exists the competitive assignment is compelling: fair in the top- n -fair and competitive senses, efficient and (essentially) unique both in terms of congestion and welfare (Proposition 3).

The definition of competitive assignments when congestion is weighted (subsection 5.2) is entirely similar (Definition 4); they are also fair, efficient and unique under minor restrictions on individual preferences (Definition 5 and Proposition 4).

In both models even with just two posts it is easy to give examples where competitive assignments do not exist. To turn Competitiveness into an operational concept of Ex Post fairness, we randomise the selection of the assignment. We assume that the agents have von Neuman Morgenstern (vNM) cardinal utilities u_i over allocations, just as in Hylland and Zeckhauser ([20]) for the non congested assignment problem.³

Section 6 focuses on the anonymous congestion model. A randomized assignment generates a *fractional* expected congestion σ_a at post a . Agent i 's competitive demand at the congestion profile $\sigma = (\sigma_x)_{x \in A}$ is the set $D(u_i, \sigma)$ of posts maximis-

²Under anonymous congestion n is also $\frac{1}{m}$ -th of the number of allocations.

³However we cannot as they did define a competitive price in fiat money because the worth of a post depends heavily upon its consumption by other agents.

ing $u_i(x, \sigma_x)$ over A . The profile σ is competitive if it obtains as a (finite) convex combination of deterministic assignments where each agent i is always assigned to a post in $D(u_i, \sigma)$ (Definition 6). The key fact is that every problem (A, N, u) has a *unique competitive congestion profile* σ^c (Lemma 1). The profile σ^c is implemented⁴ by a lottery over the deterministic assignments described above, and satisfying two properties: each agent is always assigned to a post in their competitive demand; the deterministic congestion s_a at each post a rounds the competitive congestion σ_a up or down to the nearest integer (Definition 7). Lemma 2 says that the competitive congestion profile can be so implemented, and Theorem 1 (our second main result together with Propositions 1,2) explains that each corresponding deterministic assignment is approximately top- n -fair, competitive and efficient, and that they all have approximately identical utilities.

We can randomise a weighted congestion problem in exactly the same way, and postpone this discussion to the Appendix (subsection 8.3) because the results are weaker. The fractional competitive congestion exists and is still unique under some qualifications (Definition 8 and lemma 3). But the deterministic assignments implementing it no longer approximate the competitive congestion because the relative difference between individual weights is unbounded.

The remark at the beginning of section 6 explains why the popular randomisation of the non congested assignment problem by simply extending each ordinal preference ordering to its (incomplete) stochastic dominance ordering of random allocations is in fact not useful under congestion.

Concluding comments in section 7 summarise our findings and suggests further research directions: by adding upper and lower limits to the congestion of each post; and by relaxing the strictly negative perception of congestion. The Appendix contains a couple of proofs.

1.2 related literature

1) Two independent recent papers introduce congestion in school choice. Closest to ours is Phan et al. [30] adding the congestion dimension to the school choice model ([1]): students' preferences over schools depend also on their crowding level

⁴We use this term in a non strategic sense: it means that σ^c is a convex combination of finitely many deterministic assignments,

measured by the per capita resources at each school. This is formally isomorphic to our model. To manage crowding the paper adopts a market clearing definition of ex post fairness adapting envy-freeness like we do. But the school choice viewpoint adds hard capacity constraints as well as priorities over students for the schools: this qualifies envy-freeness as “no justified envy” and their novel concept of Rationing Crowding Equilibrium (RCE) is more complex than our competitive assignment; also the proof that a RCE exists is significantly more involved. Remarkably the RCEs have the semi-lattice structure from which emerges a maximal RCE student optimal and unique welfarewise, as in the standard school choice model.

Phan et al’s model is deterministic and relies on ordinal preferences, so it cannot take advantage of the convexification offered by randomisation and vNM cardinal utilities. Still, up to a rounding argument (reminiscent of the one in our Lemma 2) in the definition of RCEs, Theorem 2 showing that the RCEs have the same crowding profile is analog to statement *i*) in our Proposition 3.

Copland [13] also defines a deterministic school choice model where students have strict ordinal preferences over allocations (a, s) ; instead of envy-freeness he imposes what we call the *free mobility* equilibrium property (on which more below) to balance fairly the congestion across schools. This property implies our top- n -fairness but is substantially more permissive than envy-freeness. Copland then proposes a variant of the Deferred Acceptance mechanism (DA with Voluntary Withdrawals) to take into account the schools’ priorities and capacity constraints. So the formal similarities with our results are fewer than in [30] but the general viewpoint on congestion is still the same.

The recent follow up by Chen et al. [10] show that deciding whether a deterministic competitive assignment (Definition 3 subsection 5.1) exists can be done in polynomial time w. r. t. m and n , whereas deciding whether an Envy-Free and top-fair assignment exists is NP-complete: in other words Competitiveness can easily be applied in large congestion problems.

3) As mentioned in the introduction, an important result for congestion games apply to precisely our model with preferences decreasing in *anonymous* congestion: the non cooperative game where agents choose their posts independently always has one or more Nash equilibria ([25], [23]).⁵ Similar games play a key role in the

⁵This is not true any more if congestion is weighted: [25].

hedonic model of coalition formation ([2], [5]) and local public goods (e. g. [6]).

The Free Mobility (FM) equilibria are logically squeezed between our two fairness concepts: every FM equilibrium is top-fair and every competitive assignment is a FM equilibrium (the converse statements do not hold). But multiple FM equilibria with different welfare consequences and no clear selection is a common situation, therefore the strategic analysis does not help us find a normatively compelling single-valued fair assignment. See section 3 and subsection 5.1 for some examples.

4) The classic combinatorial optimisation problems known as bin packing or knapsack ([11]) discuss like us how to fill bins (posts) with indivisible balls (agents); but the balls there are just objects and the concern is about the welfare of the bins (e.g., to respect a capacity constraint or to minimise the load) while we focus on the welfare of the balls and treat the bins as objects.

One exception is [12] where each ball has its own maximal acceptable congestion in each bin: this is exactly like our description of top- n - guarantees (section 2.2), with the difference that the caps are exogenous in [12] so that it may not be feasible to assign all balls. The paper shows the complexity of computing the maximal number of balls we can assign and evaluates the price of anarchy of the corresponding free mobility equilibrium.

2 the ordinal model: weighted and unweighted

We have m posts denoted a, b, \dots , and their set is A . Each agent i in the finite set N of cardinality n must be assigned to some post a in A .

An *assignment* of agents to posts is a (quasi) partition $P = (S_a)_{a \in A}$ of N where S_a is the set of agents assigned to post a ; the sets S_a, S_b are mutually disjoint and at most $m - 1$ of them can be empty. The set of all assignments is $\mathcal{P}(A, N)$, or simply \mathcal{P} .

anonymous congestion: the unweighted model Each agent adds one unit of congestion at a . Given an assignment $P \in \mathcal{P}(A, N)$ the congestion at post a is $s_a = |S_a|$ (the cardinality of S_a); we call $s = (s_a)_{a \in A}$ the *congestion profile* of P .

With the notation $[q]$ for the interval $\{1, \dots, q\}$ in \mathbb{N} , agent i 's (transitive and complete) preference relation \succeq_i bears on the set $\mathcal{F} = A \times [n]$ of feasible allocations (a, s_a) : i is assigned to post a where the congestion is s_a . Preferences are strictly

decreasing in s_a and otherwise arbitrary. The choice of A, N and a preference profile $\succeq = (\succeq_i)_{i \in N}$ defines a congestion problem (A, N, \succeq) .

agent-specific congestion: the weighted model Each agent i brings the amount w_i of congestion to her assigned post: w_i is a *strictly positive* real number and the total congestion is $W = \sum_{i \in N} w_i$. For any subset S of N we use the notation $w_S = \sum_{j \in S} w_j$.

Given an assignment $P \in \mathcal{P}(A, N)$, the congestion profile is now $s = (w_{S_a})_{a \in A}$. The set of feasible allocations (a, s_a) becomes agent-specific: $\mathcal{F}_i = A \times [w_i, W]$; agent i has continuous preferences \succeq_i over \mathcal{F}_i strictly decreasing in the congestion coordinate. A congestion problem is a 4-tuple (A, N, w, \succeq) .

The discussion of both models with ordinal preferences is the subject of sections 3 to 5.

3 the canonical guarantee: anonymous congestion

Ex ante fairness provides to each agent i a worst case welfare level that only depends, besides their own preferences, upon the number of other participants: our agent is essentially guaranteed one of their n best out of the $m \times n$ allocations in \mathcal{F} .

Fixing agent i with preferences \succeq_i on \mathcal{F} a *n-prefix* of \succeq_i is a subset of n allocations such that each one is weakly preferred to all $(m - 1)n$ other allocations. It is unique if \succeq_i is strict. Welfare decreases strictly in congestion therefore if $z = (a, s)$ is in a certain n -prefix, so does (a, s') for all $s' \leq s - 1$. So we can represent a n -prefix of \succeq_i by a vector $\lambda_i = (\lambda_{ia})_{a \in A}$ such that

$$\lambda_{ia} \in \mathbb{N} \cup \{0\} \text{ for all } a, \text{ and } \sum_{a \in A} \lambda_{ia} = n \quad (1)$$

The n -prefix is the union over all a of the sets $\{a\} \times [\lambda_{ia}]$ in \mathcal{F} , with the convention $\{a\} \times [0] = \emptyset$. We write $\Delta^{\mathbb{N}}(A; n)$ for the set of n -prefixes defined by (1).

For instance the n -prefix $\lambda_{ia} = n, \lambda_{ib} = 0$ for $b \neq a$ means that i (weakly) prefers post a to any other post x , no matter how congested x and a are. Another example has $\lambda_{ia} = \lfloor \frac{n}{m} \rfloor$ or $\lceil \frac{n}{m} \rceil$ for all a (where $\lfloor x \rfloor$ and $\lceil x \rceil$ are the smallest and largest integer bounded above and below by x): this agent's priority is to minimise the

congestion level at whatever post they are assigned (up to the necessary rounding up or down).

When indifferences in the preference \succeq_i allow several n -prefixes λ_i , it is easy to check that any two of these differ by at most 1 in each coordinate λ_{ia} .

Under anonymous congestion we define Ex Ante fairness as the guarantee that each agent's allocation is in one of their n -prefixes.

Definition 1 *Given a problem (A, N, \succeq) and a profile $\lambda = (\lambda_i)_{i \in N}$ of corresponding n -prefixes (one $\lambda_i \in \Delta^{\mathbb{N}}(A; n)$ for each preferences \succeq_i) the assignment $P \in \mathcal{P}$ is top- n -fair iff*

$$s_a \leq \lambda_{ia} \text{ for all } a \in A \text{ and } i \in S_a \quad (2)$$

We write $\mathcal{P}(\lambda)$ for the set of top- n -fair assignments and $\mathcal{C}(\lambda)$ for the set of congestion profiles s when P varies in $\mathcal{P}(\lambda)$.

Proposition 1 *In the problem (A, N, \succeq) there exists at least one top- n -fair assignment P for any profile λ of corresponding n -prefixes.*

Proof We use a simple greedy algorithm and an induction argument on n .

If for some post a we have $\lambda_{ia} = 0$ for all i we set $S_a = \emptyset$ and it remains to prove the claim on the residual problem $(A \setminus \{a\}, N, \lambda)$. We clean up in this way all posts that nobody accepts and we assume from now on that $\max_i \lambda_{ia} \geq 1$ for all a .

Clearly if the result holds for a given n , the existence of a top- n -fair assignment holds as well for all profiles $\tilde{\lambda}$ weakly larger than $\lambda \in \Delta^{\mathbb{N}}(A; n)$ in all coordinates.

Pick any post a and order the caps $\lambda_{ia}, i \in N$ as $\lambda^{*1} \geq \lambda^{*2} \geq \dots \geq \lambda^{*n}$. Write \tilde{k} for the largest k s.t. $\lambda^{*k} \geq k$ (well defined because $\lambda^{*1} \geq 1$) and pick a \tilde{k} -prefix S_a of the following ordering of N : $i \dashv^a j \Leftrightarrow \lambda_{ia} \geq \lambda_{ja}$ of N . Then $\{\lambda_{ia}\}_{i \in S_a} = \{\lambda^{*k}\}_{1 \leq k \leq \tilde{k}}$ and $\lambda_{ja} \leq \tilde{k}$ for each $j \in N \setminus S_a$ by definition of \tilde{k} .

Assigning S_a to a meets inequalities (2) for a and in the residual problem $(A \setminus \{a\}, N \setminus S_a, \lambda)$ we have $\sum_{b \in A \setminus \{a\}} \lambda_{jb} \geq n - \tilde{k} = |N \setminus S_a|$ for all j . The induction assumption on n concludes the proof. ■

Importantly the top- n -guarantee is *maximal*: if all preferences \succeq_i coincide then in any assignment P and for any common n -prefix λ_0 , at least one agent gets at P their least preferred allocation in λ_0 : therefore the top- n -guarantee cannot be improved.

We illustrate the power of top- n -fairness (2) in three examples. In the first two $\mathcal{C}(\lambda)$ is a singleton and we show the simple way to recognise this fact.

Example 1: two posts. We check that there is a single top- n -fair congestion profile: $|\mathcal{C}(\lambda)| = 1$. Set $A = \{a, b\}$ and pick a profile $\lambda = (\lambda_i)_{i \in N}$ of n -prefixes. Label the agents from 1 to n so that λ_{ia} decreases (weakly) with i , while λ_{ib} increases (weakly). Keeping in mind $\lambda_{ia} + \lambda_{ib} = n$, the integer \tilde{k} in the proof of Lemma 1 is defined by the inequalities $\lambda_{\tilde{k}a} \geq \tilde{k} \geq \lambda_{(\tilde{k}+1)a} \iff \lambda_{\tilde{k}b} \leq n - \tilde{k} \leq \lambda_{(\tilde{k}+1)b}$.

We see that the maximal congestion at a compatible with top- n -fairness is \tilde{k} while at b it is $n - \tilde{k}$: therefore $(\tilde{k}, n - \tilde{k})$ is the only top- n -fair congestion. Moreover if $\lambda_{\tilde{k}} \neq \lambda_{(\tilde{k}+1)}$ there is also a single top- n -fair assignment.

Example 2: Here $m = 4, n = 18$, and there is a single top- n -fair assignment: $|\mathcal{P}(\lambda)| = 1$.

Agents are of five types labeled α to ε . Agents of a given type have the same unique 18-prefix but not necessarily identical preferences, and the profile λ is

	a	b	c	d
$\alpha\alpha\alpha$	3	2	10	3
$\beta\beta\beta\beta$	7	1	7	3
$\gamma\gamma$	3	8	3	4
$\delta\delta\delta\delta\delta$	4	2	9	3
$\varepsilon\varepsilon\varepsilon\varepsilon$	4	0	7	7

Check that $\mathcal{P}(\lambda)$ contains just one assignment P

$$P: \begin{array}{cccc} a & b & c & d \\ \beta\beta\beta\beta & \gamma\gamma & \alpha\alpha\alpha\delta\delta\delta\delta & \varepsilon\varepsilon\varepsilon\varepsilon \end{array} \quad (3)$$

We can fit top-18-fairly at most 4 agents at post a ; at most 2 at post b ; at most 8 at post c ; at most 4 at post d . As $4 + 2 + 8 + 4 = 18$ the only top-18-fair congestion profile is $(4, 2, 8, 4)$; next we see that the only way to fit 8 agents at c is with all α -s and δ -s; then we must assign the γ -s to b and $\mathcal{P}(\lambda) = \{P\}$ follows.

More generally write $c^{mx}(a; \lambda)$ for the maximal number of agents we can fit λ -fairly at post a . Lemma 1 implies $\sum_{a \in A} c^{mx}(a; \lambda) \geq n$ for any λ . If $\sum_{a \in A} c^{mx}(a; \lambda) = n$ then all top- n -fair assignments have the same congestion profile $s_a = c^{mx}(a; \lambda)$ for all a . The converse property holds as well (we omit the easy proof):

$$|\mathcal{C}(\lambda)| = 1 \iff \sum_{a \in A} c^{mx}(a; \lambda) = n \quad (4)$$

Examples where $|\mathcal{P}(\lambda)| > |\mathcal{C}(\lambda)| = 1$ abound, e. g. Example 1 if $\lambda_{\tilde{k}} \neq \lambda_{(\tilde{k}+1)}$, or our next example.

Example 3: Here $m = 3, n = 12$ and $|\mathcal{C}(\lambda)| = 15, |\mathcal{P}(\lambda)| = 953$.

Agents are of two types, six in each type with identical 12-prefixes:

$$\begin{array}{rcc} & a & b & c \\ \alpha\alpha\alpha\alpha\alpha\alpha & 6 & 4 & 2 \\ \beta\beta\beta\beta\beta\beta & 2 & 4 & 6 \end{array}$$

There are three top-12-fair congestion profiles respecting the symmetry between the two types:

$$\begin{array}{rcccl} & a & b & c & \\ P_1 & \alpha\alpha\alpha\alpha\alpha\alpha & 0 & \beta\beta\beta\beta\beta\beta & \\ P_2 & \alpha\alpha\alpha\alpha\alpha & \alpha\beta & \beta\beta\beta\beta\beta & \\ P_3 & \alpha\alpha\alpha\alpha & \alpha\alpha\beta\beta & \beta\beta\beta\beta & \end{array} \quad (5)$$

In addition $\mathcal{C}(\lambda)$ contains six asymmetric congestion profiles $(s_a, s_b, s_c) = (6, 1, 5), (6, 2, 4), (6, 3, 3), (6, 4, 2), (5, 4, 3), (5, 3, 4)$, and six more by exchanging the role of α and β .

Count next the top- n -fair assignments: in (5) P_1 allows just one assignment, but P_2 and P_3 allow respectively 36 and 225 assignments by permuting agents within their types.⁶ Taking all profiles of $\mathcal{C}(\lambda)$ into account gives $|\mathcal{P}(\lambda)| = 953$.

Depending on individual preferences, some of the top-12-fair assignments (but not all⁷) may be inefficient (Pareto dominated). For instance if in P_3 four α -agents prefer $(a, 4)$ to $(b, 4)$ and the last two have the opposite preference then only one of the 225 corresponding assignments of the α -s between posts a and b is efficient.

two strategic interpretations of top- n -fairness Consider the direct revelation mechanism where agents report independently their preferences and the mechanism implements a reportedly top- n -fair assignment. Then every Nash equilibrium of such mechanism must be top- n -fair because the truthful report ensures this. The same is true in any mechanism in which the last stage offers to each agent the chance to claim a top- n -fair allocation.

⁶For P_3 at post b fifteen choices of the β -s and fifteen of the α -s .

⁷If $P \in \mathcal{P}(\lambda)$ is efficient within $\mathcal{P}(\lambda)$, it is clearly efficient within the entire feasible set.

For the second interpretation we suppose that agent i , clueless about other agents' preferences at the interim stage, wants to maximise their worst case welfare: reporting a truthful n -prefix of \succeq_i ensures this and no other report does. Indeed in the adversarial case where everyone else reports the same λ_i an assignment rule treating equals as equally as possible can give agent i anyone of the allocations (a, λ_{ia}) , in particular a worst one in this set. Of course, any information about other participants' preferences can be used to advantage by our agent but at some risk.

free mobility equilibria and top- n -fairness In the FM normal form game each agent picks their post a in A and consumes (a, s_a) where $s_a - 1$ is the number of other agents who chose a . We check that every Nash equilibrium is a top- n -fair assignment. Fix a profile of strategies $x = (x_i)_{i \in N} \in A^N$ and write $s(a|x)$ for the resulting congestion at post a . The equilibrium property for i is $(x_i, s(x_i|x)) \succeq_i (a, s(a|x) + 1)$ for all $a, a \neq x_i$: at such post a the only allocations i may strictly prefer to $(x_i, s(x_i|x))$ are in $\{a\} \times [s(a|x)]$ so their number is at most $s(a|x)$; adding to those the allocations in $\{x_i\} \times [s(x_i|x) - 1]$ that our agent definitely prefers to $(x_i, s(x_i|x))$, we see that at most $(n - 1)$ allocations can beat $(x_i, s(x_i|x))$ and conclude that the latter belongs to at least one n -prefix of agent i .

Recall that the FM game has at least one Nash equilibrium assignment ([25]). Clearly these assignments can be a strict subset of $\mathcal{P}(\lambda)$; e. g. in Example 3 P_1 is a FM eq. assignment only if $(a, 6) \succeq_\alpha (b, 1)$ and $(c, 6) \succeq_\beta (b, 1)$.

4 the canonical guarantee: weighted congestion

In the interim stage at which we define Ex Ante fairness, agent i knows their own weight w_i and the total congestion W , but neither the number of other agents filling the congestion $W - w_i$ or their individual weights. The welfare level guaranteed to i depends only upon these two variables and i 's preferences. This definition does not directly generalise that of guarantees under anonymous congestion because adversarial situations allows for any finite split of the total weight $W - w_i$ of other participants.

The (Lebesgue measure) size of the set of agent i 's allocations $\mathcal{F}_i = A \times [w_i, W]$ is $m(W - w_i)$: we show below that agent i can guarantee one of their top $\frac{1}{m}$

allocations in \mathcal{F}_i , i. e., their best subset of size $(W - w_i)$. Exactly like under anonymous congestion where out of $m \times n$ allocations in \mathcal{F} , agent i 's n -prefixes are the top $\frac{1}{m}$ -th quantiles of i 's preferences.

The description of agent i 's $(W - w_i)$ -prefix is more subtle than for n -prefixes, but on the bright side this prefix is unique for each preference \succeq_i . With the notation $[[z]]$ for the number of strictly positive coordinates of $z \in \mathbb{R}_+^A$, a $(W - w_i)$ -prefix is described by a vector $\lambda_i = (\lambda_{ia})_{a \in A}$ such that

$$\text{for each } a : \lambda_{ia} = 0 \text{ or } w_i \leq \lambda_{ia} \leq W, \text{ and } \sum_{a \in A} \lambda_{ia} = W + ([[\lambda_i]] - 1)w_i \quad (6)$$

The $(W - w_i)$ -prefix captured by the vector λ_i above is as follows. If $\lambda_{ia} \in [w_i, W]$ it contains all allocations in $\{a\} \times [w_i, \lambda_{ia}]$. In particular $\lambda_{ia} = w_i$ means that agent i can be at post a only if she is alone there. If $\lambda_{ia} = 0$ agent i cannot be assigned to post a .

Notice that changing λ_{ia} from 0 to w_i adds w_i to both sides of (6), therefore ruling out post a out or accepting it only if i is alone there has no impact on the equality constraint in (6). For instance if $\lambda_{ia} = W$ our agent can still choose to accept any subset of the other posts provided she is alone there.

Let B be the (strict and possibly empty) subset B of the posts such that $\lambda_i = 0$; the total size of the allocations in $(A \setminus B) \times [w_i, W]$ that λ_i allows is

$$\sum_{a \in A \setminus B} (\lambda_{ia} - w_i) = \sum_{a \in A} \lambda_{ia} - (m - |B|)w_i = W - w_i$$

therefore λ_i cuts a subset of size $(W - w_i)$ in \mathcal{F}_i as announced two paragraphs before (6). Moreover for any preference \succeq_i the top $\frac{1}{m}$ quantile in $\mathcal{F}_i = A \times [w_i, W]$ is unique: this is clear by continuity of preferences and strictness in each $\{a\} \times [w_i, W]$. Then we set $\lambda_{ia} = 0$ if and only if this set does not contain any allocation at post a and λ_i is uniquely defined as well.

Definition 2 *In the problem (A, N, \succeq, w) where λ_i meeting (6) is the $(W - w_i)$ -prefix of \succeq_i , an assignment P is top- $\frac{1}{m}$ -fair if and only if*

$$\sum_{j \in S_a} w_j \leq \lambda_{ia} \text{ for all } a \in A \text{ and all } i \in S_a$$

Proposition 2 *In any problem (A, N, \succeq, w) there exists at least one top- $\frac{1}{m}$ -fair assignment P .*

Proof in the Appendix section 8.1. There we also state and prove that the canonical guarantee is maximal, just like when congestion is anonymous.

Finding the top- $\frac{1}{m}$ -fair assignments is more difficult with weighted rather than unweighted congestion. This is already clear with two posts.

Example 4: *Two posts and three agents, $m = 2, n = 3$, and $W = 10$. The weights w_i and $(W - w_i)$ -prefixes are*

	<i>a</i>	<i>b</i>	<i>w</i>
α	10	6	6
β	6	6	2
γ	8	4	2

Here α 's feasible set is $\mathcal{F}_\alpha = \{a\} \times [6, 10] \cup \{b\} \times [6, 10]$ and its 4-prefix $\{a\} \times [6, 10] \cup \{b, 6\}$; similarly γ accepts allocations in $\{a\} \times [2, 8] \cup \{b\} \times [2, 4]$.

Suppose every agent i 's preferences over allocations are represented by the *slack utilities*:⁸

$$u_i(x, w_{S_x}) = \lambda_{ix} - w_{S_x} \text{ for all } x \tag{7}$$

We find three top- $\frac{1}{2}$ -fair assignments with very different congestion and utilities

	<i>a</i>	<i>b</i>		α	β	γ
assignments:	P_1	α	$\beta\gamma$	utilities:	P_1	4 2 0
	P_2	$\alpha\gamma$	β		P_2	2 4 0
	P_3	$\beta\gamma$	α		P_3	0 2 4

Recall that with anonymous congestion and strict preferences (Example 1 section 3) two-post problems have a unique top- n -assignment.

A Free Mobility equilibrium with weighted congestion may have no equilibrium (see an example in section 8 of [25]) but if it does, the corresponding assignment is top- $\frac{1}{m}$ -fair. Fix such an assignment P where (a, w_{S_a}) is agent i 's equilibrium allocation. To (a, w_{S_a}) our agent prefers the allocations $\{a\} \times [w_i, w_{S_a}]$ and possibly those in $\{x\} \times [w_i, w_{S_x} + w_i]$: the total length of those sets is exactly $W - w_i$, which proves the claim.

⁸This simple utility function derived from just its individual prefix is frequently used in the rest of the paper, both with ordinal preferences and cardinal utilities.

5 deterministic competitive assignments

5.1 under anonymous congestion

We use the notation $z \vee y = \max\{z, y\}$.

Definition 3 *In problem (A, N, \succeq) the assignment P is competitive (Comp) iff*

$$\text{for all } a \in A \text{ and } i \in S_a : (a, s_a) \succeq_i (x, s_x \vee 1) \text{ for all } x \in A \quad (8)$$

If all posts are occupied ($s_x \geq 1$ for all x) property (8) is just envy-freeness: switching my allocation to yours is a virtual move that does not affect the congestion at your post (unlike an actual move in the FM game). Interpreting the congestion s_x at post x as the “price” of consuming x , my competitive demand is the allocation (x, s_x) when I assume that the price vector is fixed.

Property (8) assigns a price of 1 to an empty post x and insists that our agent does not want to move there either: this is important because an assignment where all agents share a single post a is automatically envy-free, but can be the absolute worst assignment for everyone.

Proposition 3 *Fix a problem (A, N, \succeq) .*

- i) All competitive assignments have the same congestion profile (except possibly at some posts occupied by at most one agent), and the same welfare profile.*
- ii) A competitive assignment is weakly efficient, and efficient if preferences are strict and/or if all posts are occupied.*
- iii) A competitive assignment is top- n -fair.*

Proof At the congestion profile $s \in \Delta^{\mathbb{N}}(A; n)$ ((1)) we define agent i 's *competitive demand* as

$$D(i; s) = \{a \mid (a, s_a \vee 1) \succeq_i (x, s_x \vee 1) \text{ for all } x \in A\} \quad (9)$$

The assignment P is competitive if and only if $a \in D(i; s)$ whenever $i \in S_a$.

Statement i) Unique congestion. We fix two different congestion profiles s, s^* coming from the competitive assignments $P = (S_x)_{x \in A}$ and $P^* = (S_x^*)_{x \in A}$.

Define the set $A^* = \{a \in A : s_a^* \vee 1 > s_a \vee 1\}$ and assume that A^* is non empty, which will lead to a contradiction. Note that in A^* we have $s_a^* > s_a, 1$.

Fixing an agent i we claim that if $D(i; s^*)$ intersects A^* then $D(i; s)$ must be a subset of A^* . If the claim fails there is some $b \in D(i; s)$ outside A^* such that for all a

in A^* : $(b, s_b \vee 1) \succeq_i (a, s_a \vee 1) \succ_i (a, s_a^*)$ (strict preference because $s_a^* > s_a, 1$). By the choice of b we also have $s_b \vee 1 \geq s_b^* \vee 1$; therefore $(b, s_b^* \vee 1) \succeq_i (b, s_b \vee 1) \succ_i (a, s_a^*)$, and as a was arbitrary in A^* it follows that $D(i; s^*)$ cannot intersect A^* , contradiction.

Now for each $i \in \cup_{a \in A^*} S_a^*$ the claim says $D(i; s) \subseteq A^*$ therefore i is assigned to A^* by P . This implies $\sum_{a \in A^*} s_a \geq \sum_{a \in A^*} s_a^*$, contradicting $s_a^* > s_a$ in A^* . We conclude that A^* is empty.

So P, P^* must be such that $s_a^* \vee 1 = s_a \vee 1$ for all a : SO $s_a^* \neq s_a$ can only happen when $\{s_a^*, s_a\} = \{0, 1\}$, as claimed in statement i). Moreover the competitive demands at s and s^* coincide.

A simple example with multiple competitive congestion profiles has all agents except 1 and 2 refusing the three posts a, b, c , while 1 and 2 refuse all but a, b, c and they are indifferent between $(a, 1), (b, 1)$ and $(c, 1)$: combining a competitive sub-assignment of $N \setminus \{1, 2\}$ to $A \setminus \{a, b, c\}$ with any assignment where agents 1, 2 occupy two of a, b, c is competitive in the full problem.

Unique welfare. We fix an agent i in S_a and S_b^* and show that i is indifferent between the two assignments. From $s_a, s_b^* \geq 1$ and Definition 3 we have

$$(a, s_a) \succeq_i (b, s_b \vee 1) \text{ and } (b, s_b^*) \succeq_i (a, s_a^* \vee 1) \quad (10)$$

If $s_a = s_a^*$ and $s_b = s_b^*$ we are done. If $s_a \neq s_a^*$ and $s_b = s_b^*$ the equality $s_a^* \vee 1 = s_a \vee 1$ implies $s_a = 1 > 0 = s_a^*$; then $s_b^* \geq 1$ and (10) give $(a, 1) \succeq_i (b, s_b) = (b, s_b^*) \succeq_i (a, 1)$ as desired. The last subcase $s_a \neq s_a^*$ and $s_b \neq s_b^*$ is just as easy.

Statement ii) Efficiency. Assume, to the contrary that $P = (S_x)$ is competitive and Pareto inferior to $Q = (T_x)$. Say agent i , assigned to a at P , is assigned to b at Q (a, b are not necessarily distinct) and suppose that post b is occupied at P : $s_b \geq 1$. Then by Comp and the weak Pareto improvement we have

$$(b, s_b) \preceq_i (a, s_a) \preceq_i (b, t_b) \implies s_b \geq t_b \quad (11)$$

and $s_b > t_b$ if agent i improves strictly at Q .

If all posts are occupied at P (11) implies $s = t$ and we have a contradiction. If instead some agent i goes from a at P to c at Q and c is empty at P , $s_c = 0$, we have

$$(c, 1) \preceq_i (a, s_a) \preceq_i (c, t_c) \implies t_c = 1 \text{ and } (a, s_a) \simeq_i (c, 1) \quad (12)$$

This is a contradiction if preferences are strict, and we conclude again that P is Pareto optimal. If indifferences are possible we see that i is a weak Pareto optimum (not all agents benefit strictly).

The argument above explains how a competitive assignment can be Pareto inferior: in the situation (12) moving i from a to c and *changing nothing else* is a Pareto improvement to a new assignment \tilde{P} where agents in $S_a \setminus \{a\}$ benefit strictly while all others are unaffected.

Note that \tilde{P} is not competitive if $s_a \geq 2$ because i prefers strictly $(a, s_a - 1)$ to (a, s_a) and therefore to $(c, 1)$; but if $s_a = 1$ then \tilde{P} is competitive and welfare-wise indifferent to P .

We have shown that Pareto improving a competitive assignment is only possible if \succeq_i is not strict and some post is unoccupied.

Statement iii) top- n -fairness. A FM equilibrium is top- n -fair (section 3). Now a competitive assignment is clearly a FM equilibrium. ■

Proposition 3 vindicates a competitive assignment as the essentially single-valued, fair and efficient solution to the congested assignment problem, precisely what we set out to discover. But it is as easy to find problems where such assignment exists as where they don't. Recall that deciding if one exists in a given problem is computable in polynomial time ([10]). Deciding when existence is "likely" in some plausible domain of preferences is therefore a feasible numerical and empirical project.

Start with an extreme example where $n = m$ and every agent's n -prefix is $\lambda_{ia} = 1$ for all a : they insist on being alone at their assigned post. An assignment is top- n -fair if and only if it is a one-to-one (non congested) assignment. Therefore a competitive one exists only if we can match each agent with one of their best posts. By contrast any assignment is a FM equilibrium outcome.

In Example 2 (section 3) P ((3)) is the the unique top-18-fair assignment. Assume that the preferences of each agent i with n -prefix λ_i are described by the *slack* utilities $u_i = \lambda_{ia} - s_a$ as in (7). Then it is easy to check that P is competitive.⁹ However, if the preferences of some α , say, change to $(b, 2) \succ_\alpha (c, 8)$ (compatible with the given n -prefixes) then P is no longer competitive.

⁹At P the congestion profile is $(a, 4), (b, 2), (c, 8), (d, 4)$ so the α -s' utility at c is 2 and at most 0 at any other post; the β -s' utility at a is 3 and at most 1 elsewhere; etc...

Turning to Example 3 (section 3) and maintaining the slack utilities assumption, we find that *none* of the hundreds of top-12-assignments is competitive: at P_1 in (5) all agents want to move to b ; at P_2 all agents at a or c envy those at b , and the reverse statement holds at P_3 . Checking all asymmetric congestion profiles is just as easy.

Our last two examples have three posts and six agents, $m = 3, n = 6$, and three pairs of identical agents with a cyclical pattern in their preferences. We use them to illustrate the welfare performance of competitive assignments within the set of FM eq. assignments (to which they belong by Definition 3).

Example 5.1 *Where the competitive assignment Pareto dominates two more FM eq. outcomes.* Each agent has a single 6-prefix, a single top- n -fair congestion profile $\mathcal{C}(\lambda) = \{(2, 2, 2)\}$ (by (4)) and ten top-6-fair assignments

$$\begin{array}{cccc}
 & a & b & c \\
 \alpha\alpha & 3 & 2 & 1 \\
 \beta\beta & 1 & 3 & 2 \\
 \gamma\gamma & 2 & 1 & 3
 \end{array}
 \quad
 \begin{array}{ccc}
 & a & b & c \\
 P_1 & \alpha\alpha & \beta\beta & \gamma\gamma \\
 P_2 & \gamma\gamma & \alpha\alpha & \beta\beta \\
 P_3 & \alpha\gamma & \alpha\beta & \beta\gamma
 \end{array}
 \tag{13}$$

because P_3 defines eight assignments by permuting agents within types.

If preferences are described by the slack utilities (7), only P_1 is competitive with utility 1 for everyone while P_2, P_3 are FM eq. outcomes Pareto inferior to P_1 .¹⁰

Example 5.2 *With a competitive assignment and other non Pareto comparable FM eq. assignments.* The unique prefixes and preferences are as follows:

$$\begin{array}{ccc}
 a & b & c \\
 \alpha\alpha & 3 & 3 & 0 \\
 \beta\beta & 3 & 0 & 3 \\
 \gamma\gamma & 0 & 3 & 3
 \end{array}
 \quad
 \begin{array}{l}
 \alpha : (a, 2) \succ_{\alpha} (b, 2) \succ_{\alpha} (b, 3) \succ_{\alpha} (a, 3) \\
 \beta : (c, 2) \simeq_{\beta} (a, 2) \\
 \gamma : (c, 1) \succ_{\gamma} (b, 2) \succ_{\gamma} (b, 3) \succ_{\gamma} (c, 2)
 \end{array}$$

The following competitive assignment P_1 is not welfare comparable to the non-

¹⁰Respectively with utility 0 for everyone at P_2 , or 0 and 1 in each type at P_3 .

competitive FM equ. assignments P_2 and P_3 .¹¹

	a	b	c
P_1	$\alpha\alpha$	$\gamma\gamma$	$\beta\beta$
P_2	$\beta\beta$	$\alpha\alpha\gamma$	γ
P_3	$\alpha\beta$	$\alpha\gamma\gamma$	β

5.2 under weighted congestion

Definition 4 In problem (A, N, \succeq, w) , the assignment P is competitive (Comp) iff

$$\text{for all } a \in A \text{ and } i \in S_a : (a, w_{S_a}) \succeq_i (x, w_{S_x} \vee w_i) \text{ for all } x \in A$$

The interpretation is the same as before, except that the ‘‘congestion price’’ of a post is now agent-specific. As before competitiveness implies top-fairness, but the critical uniqueness and efficiency properties require some qualification.

Definition 5 The assignment P is crowded iff $|S_a| \geq 2$ or $S_a = \emptyset$ for all a . The ordinal preferences \succeq_i are semi-strict iff for any distinct posts a, b and any agent i and S s. t. $i \in S \subseteq N$, agent i is not indifferent between (a, w_i) and (b, w_S) .

Proposition 4: Fix a problem (A, N, \succeq, w) .

- i)* If all preferences \succeq_i are semi-strict then all competitive assignments have the same congestion and welfare profiles. The same is true if at least one competitive assignment is crowded.
- ii)* If all preferences are semi-strict a competitive assignment is efficient.
- iii)* A competitive assignment is top- $\frac{1}{m}$ -fair

Proof in the Appendix section 8.2.

To illustrate the role of the semi-strictness assumption in statement *i)* it is enough to check that in the two-post Example 4 section 4, P_1 and P_2 are both competitive assignments with different welfare and even congestion profiles.

¹¹The α -s prefer P_1 to P_2 and P_3 (at least weakly); the β -s prefer P_3 to P_1 (at least weakly) and are indifferent between P_1 and P_2 ; the γ -s strictly prefer P_1 to P_3 and disagree when comparing P_1 to P_2 .

6 competitive fractional assignments under anonymous congestion

In the randomised version of the unweighted model we endow each agent i with a cardinal vNM utility $u_i(a, s)$ over $A \times [n]$. Fixing the post a and some expected congestion σ at a we write $u_i(a, \sigma)$ for the linear interpolation of $u_i(a, s)$ between $\lfloor \sigma \rfloor$ and $\lceil \sigma \rceil$ (the two rounded up and down values of σ):

$$u_i(a, \sigma) = \frac{\lceil \sigma \rceil - \sigma}{\lceil \sigma \rceil - \lfloor \sigma \rfloor} u_i(a, \lfloor \sigma \rfloor) + \frac{\sigma - \lfloor \sigma \rfloor}{\lceil \sigma \rceil - \lfloor \sigma \rfloor} u_i(a, \lceil \sigma \rceil) \text{ for all non-integer } \sigma \quad (14)$$

So $u_i(a, \sigma)$ is continuous and strictly decreasing in $\sigma \in [1, n]$.

The competitive approach becomes operational because it selects in every problem a unique competitive congestion profile and implements it by a lottery over approximately competitive deterministic assignments. The similar randomisation in the weighted congestion model still identifies an essentially unique congestion profile, but its implementation is less satisfactory. See subsection 8.3 in the Appendix for details.

For a finite set Z the notation $\Delta^{\mathbb{R}}(Z; y)$ is the simplex with non negative real coordinates in Z adding up to the real number y . The key variable of a random assignment is the expected congestion σ_a at each post a : the congestion profile $\sigma = (\sigma_a)_{a \in A}$ varies in $\Delta^{\mathbb{R}}(A; n)$. Note that $\sigma_a < 1$ is possible: in this case the definitions below make sure that each agent perceives a congestion of 1 at a , so $u_i(a, \sigma)$ is *not* defined if $\sigma \in [0, 1[$. Moreover in Definition 7 describing the implementation of a random congestion σ by a lottery over deterministic assignments, post a 's congestion at the latter assignments takes only the values $\lfloor \sigma_a \rfloor$ and $\lceil \sigma_a \rceil$, so that equation (14) is indeed the relevant vNM expected utility at (a, σ) .

As in the proof of Proposition 3 we define agent i 's *competitive demand* at the expected congestion profile $\sigma \in \Delta^{\mathbb{R}}(A; n)$:

$$D(u_i, \sigma) = \arg \max_{x \in A} u_i(x, \sigma_x \vee 1) \quad (15)$$

This demand ignores the effect of agent i 's own presence at post x but correctly anticipates that the ex post congestion will round σ_x up or down. Our key definition comes next.

Definition 6: In the problem (A, N, u) the fractional congestion profile $\sigma \in \Delta^{\mathbb{R}}(A; n)$ is competitive (Comp) iff

$$\sigma \in \sum_{i \in N} \Delta^{\mathbb{R}}[D(u_i, \sigma); 1] \quad (16)$$

Property (16) means that we can achieve the congestion profile σ^* by assigning randomly each agent, with a well chosen probability distribution, over the posts in her competitive demand. Writing this probability distribution as (π_{ia}) we obtain a semi-stochastic matrix $\Pi = [\pi_{ia}] \in [0, 1]^{N \times A}$ (each row sums to 1) realising σ and with competitive support:

$$\sigma_a = \sum_{i \in N} \pi_{ia} \text{ for all } a \in A \quad (17)$$

$$\{\pi_{ia} > 0 \implies a \in D(u_i, \sigma)\} \text{ for all } i \in N, a \in A \quad (18)$$

There may be more than one matrix Π realising a given congestion profile σ .

Lemma 1: In every problem (A, N, u) there is a unique competitive congestion profile σ^c , except possibly at some posts where $\sigma_a^c \leq 1$,¹² and a unique competitive demand.

Proof Existence. It follows from the Kakutani fixed point theorem applied to the convex compact valued correspondence $\Gamma(\sigma) = \sum_{i \in N} \Delta^{\mathbb{R}}(D(u_i, \sigma); 1)$ mapping $\Delta^{\mathbb{R}}(A; n)$ into itself. To check that the graph of Γ is closed (implying that Γ is upper-semi-continuous) we take a sequence $(\sigma^t, \tau^t)_{t=1}^{\infty}$ in $[\Delta^{\mathbb{R}}(A; n)]^2$ converging to (σ, τ) and s. t. $\tau^t = \sum_{i \in N} \tau_i^t$ and $\tau_i^t \in \Delta^{\mathbb{R}}[D(u_i, \sigma^t); 1]$. We can find a subsequence such that all sequences $\{\tau_i^t\}$ converge, and all sets $D(u_i, \sigma^t)$ are constant in t , so that $\tau \in \Gamma(\sigma)$ as desired.

Uniqueness. We adapt the proof of statement *i*) in Proposition 3. Assume σ, σ^* are two different competitive profiles with corresponding matrices Π and Π^* in (17), (18). We set $A^* = \{a \in A | \sigma_a^* \vee 1 > \sigma_a \vee 1\}$ and check, exactly like in the deterministic proof, that if an agent i is s. t. $D(u_i; \sigma^*)$ intersects A^* at some a , then $D(u_i; \sigma)$ must be a subset of A^* .

Therefore for any i s. t. $\sum_{a \in A^*} \pi_{ia}^* > 0$ we have $\sum_{a \in A^*} \pi_{ia}^* \leq 1 = \sum_{a \in A^*} \pi_{ia}$. Summing up over all agents this gives $\sum_{a \in A^*} \sigma_a^* = \sum_{i \in N, a \in A^*} \pi_{ia}^* \leq \sum_{i \in N, a \in A^*} \pi_{ia} =$

¹²If σ, σ^* are both competitive and $\sigma_a^* \neq \sigma_a$ at some post a , then $\sigma_a^*, \sigma_a \leq 1$.

$\sum_{a \in A^*} \sigma_a$, a contradiction of the assumption $\sigma_a^* > \sigma_a$ in A^* . We conclude that A^* is empty: $\sigma_a^* \vee 1 \leq \sigma_a \vee 1$ for all a . The opposite inequality holds by exchanging the roles of σ and σ^* , implying $\sigma_a^* \vee 1 = \sigma_a \vee 1$ and the desired conclusion that $\sigma_a^* \neq \sigma_a$ can only happen when both are at most 1.

Because an agent evaluates being assigned to a post with expected congestion at most 1 as being alone there (definition (15)) this implies $D(u_i, \sigma) = D(u_i, \sigma^*)$, therefore σ and σ^* generate the same competitive demand. ■

Definition 7: *In the problem (A, N, u) the list $(\{P^k\}_{k=1}^K, \mathcal{L})$ of K deterministic assignments P^k together with a lottery \mathcal{L} over $[K]$, implements the competitive congestion profile σ^c if, first, the expected congestion over these K assignments is σ^c : $\mathbb{E}_{\mathcal{L}}(s_a^k) = \sigma_a^c$ for all a ; and second for all $k \in [K]$, all i , all a and all k we have:*

$$i \in S_a^k \implies a \in D(u_i, \sigma^c) \text{ and } s_a^k = \lfloor \sigma_a^c \rfloor \text{ or } \lceil \sigma_a^c \rceil \quad (19)$$

Property (19) says that anyone of the assignments P^k is ex post fair in the sense that each agent is at a post in his competitive demand (based on the expected congestion σ^c) and the actual congestion at P^k is an integer rounding of the latter.

Lemma 2: *In every problem (A, N, u) we can implement the competitive congestion profile σ^c by (typically several) lists $(\{P^k\}_{k=1}^K, \mathcal{L})$ as in Definition 4.*

Corollary: *If σ^c is integer-valued each assignment P^k is competitive and implements σ^c .*

Proof Define the set \mathcal{S} of semi-stochastic matrices Π s.t. for all a and i : $\{\pi_{ia} > 0 \implies a \in D(u_i, \sigma^c)\}$ and $\lfloor \sigma_a^c \rfloor \leq \sum_{i \in N} \pi_{ia} \leq \lceil \sigma_a^c \rceil$. This set is a convex compact polytope, non empty as it contains any matrix Π^c realising σ^c . We claim that each extreme point Π^k of \mathcal{S} is deterministic, i. e., its entries are all integers or zero; so Π^k is a deterministic assignment P^k meeting (19) and Π^c is a convex combination of such extreme points: this gives us the desired collection P^k and lottery \mathcal{L} .

We prove the claim by contradiction. Pick an extreme point Π of \mathcal{S} and associate to Π the bipartite graph G on $N \times A$ containing the edge ia iff $\pi_{ia} > 0$. Extract from G the subgraph G_0 of its fractional entries ia , i. e., $0 < \pi_{ia} < 1$, and let F be the set of posts a s.t. $\sum_{i \in N} \pi_{ia}$ is fractional (not an integer or zero). If F is non empty it contains some post a : at least one edge ia is in G_0 ; then at least one other edge ib is in G_0 (Π is semi-stochastic at i); if $b \in F$ then we can add a small ε to

π_{ia} and take away ε from π_{ib} without leaving \mathcal{S} , or vice versa: this contradicts the extremality of Π , so b is not in F after all. But then there is another edge jb in G_0 because $\sum_{i \in N} \pi_{ib}$ is an integer, and again there is some new edge jc in G_0 ; if $c = a$ or $c \in F$ we can as above perturb a little the entries of Π in two opposite ways and reach a contradiction. This construction must stop at a new post in F or cycle back to a : in both cases we can perturb the path or cycle in opposite directions. The claim is proved and the proof is complete. ■

Our main result refines the properties of the deterministic assignments P^k selected ex post to implement the competitive congestion: any two such assignments yield approximately identical utilities; and each P^k shares (approximately) the properties of deterministic competitive assignments in Proposition 1. We use for agent i 's approximation parameter her worst utility loss from one additional unit of congestion $\delta_i = \max_{(a,s) \in A \times [n]} \{u_i(a,s) - u_i(a,s+1)\}$. For instance $\delta_i = 1$ for a slack utility $u_i = \lambda_{ia} - s_a$.

Theorem 1 *Fix a problem (A, N, u) , a list $(\{P^k\}_{k=1}^K, \mathcal{L})$ implementing the competitive congestion σ^c (Definition 6) and write $U^k = (U_i^k)_{i \in N}$ the utility profile at assignment P^k .*

- i) Each P^k is top- n -fair up to at most one unit of congestion: for all $i \in N$ there is a n -prefix λ_i of u_i such that $i \in S_a^k \implies s_a^k \leq \lambda_{ia} + 1$.*
- ii) Utilities at two assignments P^k, P^ℓ differ by at most $2\delta_i$: $|U_i^k - U_i^\ell| \leq 2\delta_i$ for all i and all $k, \ell \in [K]$.*
- iii) Each P^k is $2\delta_i$ -competitive: $U_i^k \geq u_i(x, s_x^k \vee 1) - 2\delta_i$ for all i and all x .*
- iv) Each assignment P^k is $(2\delta_i + \varepsilon)$ -efficient for any $\varepsilon > 0$: no assignment $Q \in \mathcal{P}(A, N)$ Pareto dominates the profile $(U_i^k + 2\delta_i + \varepsilon)_{i \in N}$.*

If in addition $\sigma_a^c \geq 1$ for all a , we can say more. Let U_i^c be the value $u_i(a, \sigma_a^c)$ common to all $a \in D(u_i, \sigma^c)$, then the profile $(U_i^c)_{i \in N}$ is efficient¹³; finally we have $U_i^k > U_i^c - \delta_i$ for all i and all $k \in [K]$.

Proof *Statement i)* Fix P^k and write its congestion profile simply as $(s_x)_{x \in A}$, then fix an agent i and her allocation (a, s_a) at P^k ; finally $[\sigma^c]$ is the support of the competitive congestion, containing x iff $\sigma_x^c > 0$.

Suppose first $s_a \geq 2$. By (19) we have $s_a - 1 \leq \lfloor \sigma_a^c \rfloor$ therefore $\sigma_a^c \geq 1$. We apply repeatedly the monotonicity of $u_i(x, s)$ in s and property (19).

¹³Not dominated by the utility profile of any assignment $Q \in \mathcal{P}(A, N)$.

For all $x \in [\sigma^c]$: $u_i(a, s_a - 1) \geq u_i(a, \lfloor \sigma_a^c \rfloor) \geq u_i(a, \sigma_a^c) \geq u_i(x, \sigma_x^c \vee 1) \geq u_i(x, \lceil \sigma_x^c \rceil)$, where the last inequality follows from $\sigma_x^c > 0$. These inequalities imply that agent i prefers to $(a, s_a - 1)$: at most $\lfloor \sigma_a^c \rfloor - 1$ less congested allocations at post a ; at most $\lceil \sigma_x^c \rceil - 1 \leq \lfloor \sigma_x^c \rfloor$ allocations at post x if $\sigma_x^c > 0$; and if $\sigma_x^c = 0$ (19) gives $u_i(a, s_a) \geq u_i(y, 1)$ so no allocation at y improves (a, s_a) or $(a, s_a - 1)$ for our agent i . We see that the number of allocations improving $(a, s_a - 1)$ is at most $(\sum_{x \in A} \lfloor \sigma_x^c \rfloor) - 1 \leq \sum_{x \in A} \sigma_x^c - 1 = n - 1$. As in Proposition 1 we conclude that $(a, s_a - 1)$ is top- n -fair, and we are done.

In the case $s_a = 1$ the argument is simpler. The allocation $(a, 1)$ is the best at post a ; at any other post x the inequality $u_i(a, 1) \geq u_i(x, \sigma_x^c \vee 1)$ implies it can be improved by at most $\lfloor \sigma_x^c \rfloor - 1$ less congested allocations (or zero if $\sigma_x^c < 1$). Then the allocation (a, s_a) itself is top- n -fair.

Statement ii) Fix P^k, P^ℓ, a, b and $i \in S_a^k \cap S_b^\ell$. By (19) a, b are both in $D(u_i, \sigma^c)$ hence $u_i(a, \sigma_a^c \vee 1) = u_i(b, \sigma_b^c \vee 1)$. By (19) again and the definition of δ_i we have $|u_i(a, s_a^k) - u_i(a, \sigma_a^c \vee 1)| \leq \delta_i$ and $|u_i(b, s_b^\ell) - u_i(b, \sigma_b^c \vee 1)| \leq \delta_i$, so the desired inequality follows.

Statement iii) We fix $i \in S_a^k$ and x then combine three inequalities: $a \in D(u_i, \sigma^c)$ gives $u_i(a, \sigma_a^c \vee 1) \geq u_i(x, \sigma_x^c \vee 1)$; next $|u_i(a, s_a^k) - u_i(a, \sigma_a^c \vee 1)| \leq \delta_i$ follows from $|s_a^k - \sigma_a^c| < 1$ ((19)) and the definition of δ_i ; similarly $|u_i(x, s_x^k) - u_i(x, \sigma_x^c \vee 1)| \leq \delta_i$.

Statement iv) For the first part of the statement we fix P^k and $Q = (T_x)_{x \in A}$ in $\mathcal{P}(A, N)$. There is at least one post b in the support of Q ($T_b \neq \emptyset$) such that $s_b^k \leq t_b$: otherwise $t_b < s_b^k$ for each b s. t. $t_b \geq 1$, which contradicts $\sum_A t_b = \sum_A s_b^k = n$. Pick such a post b , some $i \in T_b$, and let a be the post assigned to i by P^k ($a = b$ is possible). By statement *iii)* we have $u_i(a, s_a^k) \geq u_i(b, s_b^k \vee 1) - 2\delta_i$ and $u_i(b, s_b^k \vee 1) \geq u_i(b, t_b)$ by our choice of b . So Q does not improve P^k by more than $2\delta_i$ for all agents in T_b .

For the second part we have $\sigma_a^c \geq 1$ for all a . Assume that the utility profile of a deterministic assignment $Q = (T_x)_{x \in A}$ Pareto dominates $(U_i^c)_{i \in N}$. For any post a in the support of Q and agent $i \in T_a$, the definition of the competitive demand ((15)) and the choice of Q imply $u_i(a, \sigma_a^c) \leq U_i^c \leq u_i(a, t_a)$ and in turn $t_a \leq \sigma_a^c$ for all a . As $\sum_a t_a = \sum_a \sigma_a^c$ we conclude $t = \sigma^c$ and the inequalities above are equalities, so Q 's utility profile is exactly $(U_i^c)_{i \in N}$, and Q does not dominate $(U_i^c)_{i \in N}$ after all.

Next for any P^k , any i and a such that $i \in S_a^k$, Definition 4 implies $s_a^k \leq \lceil \sigma_a^c \rceil <$

$\sigma_a^c + 1$ hence $u_i(a, s_a^k) > u_i(a, \sigma_a^c + 1) \geq u_i(a, \sigma_a^c) - \delta_i = U_i^c - \delta_i$ as desired. ■

The simplest examples of a fractional competitive assignments have only two posts.

Example 6 *Two posts and eight agents of three types, $m = 2, n = 8$. The prefixes allow a single top- n -fair assignment*

$$\begin{array}{rcc} & a & b \\ \alpha\alpha\alpha & 8 & 0 \\ \beta\beta & 4 & 4 \\ \gamma\gamma & 0 & 8 \end{array} \quad \text{top-8-fair } P_1 : \begin{array}{rcc} & a & b \\ \alpha\alpha\alpha & \beta\beta\gamma\gamma & \end{array}$$

P_1 is competitive if and only if both β -s prefer (at least weakly $(b, 4)$ to $(a, 4)$). We assume instead $(a, 4) \succ_\beta (b, 4)$ for both β -s, so at P_1 they both envy post a . Their cardinal utilities differ slightly: β_1 suffers comparatively less than β_2 at $(a, 5)$

$$\begin{array}{rcccc} & (a, 5) & (b, 4) & (a, 4) & (b, 3) \\ u_{\beta_1} & 0 & 0 & 1 & 2 \\ u_{\beta_2} & -1 & 0 & 1 & 2 \end{array}$$

The competitive fractional congestion must load post a more than post b : $\sigma^c = (4+x, 4-x)$ for some $x \geq 0$, therefore some deterministic assignment implementing σ^c will violate top-8-fairness. By statement *i*) in Theorem 1 we have $x \leq 1$, so the demand of the α -s and γ -s will not change. To meet property (16) we need $D(u_{\beta_i}, \sigma) = \{a, b\}$ for $i = 1$ or 2 : the correct choice is $D(u_{\beta_2}, \sigma) = \{a, b\}$, implying $x = \frac{1}{4}$ then $D(u_{\beta_1}, \sigma) = \{a\}$.¹⁴

The lottery $\frac{3}{4}P_1 + \frac{1}{4}P_2$ where

$$P_2 = \begin{array}{rcc} & a & b \\ \alpha\alpha\alpha\alpha\beta_2 & \beta_1\gamma\gamma & \end{array}$$

implements $\sigma^c = (4\frac{1}{4}, 3\frac{3}{4})$ as specified by Theorem 1.

Note that for agent β_2 , conditional on being at post a the expected congestion there is 5, not $4\frac{1}{4}$: reasoning competitively β_2 accepts both posts at σ^c without taking into account the impact of their own assignment to a on the congestion there. By property (19) this discrepancy is at most one unit of congestion, irrespective of the size of n and m .

¹⁴Choosing $D(u_{\beta_1}, \sigma) = \{a, b\}$ forces $x' = \frac{1}{3}$ then $D(u_{\beta_2}, \sigma) = \{b\}$, and (16) fails.

Comparing the competitive lottery $\frac{3}{4}P_1 + \frac{1}{4}P_2$ and the top- n -fair assignment P_1 (also the only FM equ. outcome), agent β_1 and the γ -s prefer the lottery while β_2 and the α -s have the opposite preference.

Example 3 continued from section 3. $m = 3, n = 12$.

The two types of agents have the following 12-prefixes

	a	b	c
$\alpha\alpha\alpha\alpha\alpha$	6	4	2
$\beta\beta\beta\beta\beta$	2	4	6

Suppose ordinal preferences are captured by the slack utilities (7). Their vNM extension to real valued congestion (14) is still linear in congestion: $u_i(a, \sigma_a) = \lambda_{ia} - \sigma_a$.

We checked in section 3 that no deterministic assignment is competitive. The unique competitive fractional congestion respects the symmetries of the problem (this is a general property following its definition by the fixed point property (16)), so we look for $\sigma^c = (x, y, x)$ with $2x + y = 12$. The only choice generating the demands $D(u_\alpha, \sigma) = \{a, b\}$, $D(u_\beta, \sigma) = \{b, c\}$ is $\sigma^c = (4\frac{2}{3}, 2\frac{2}{3}, 4\frac{2}{3})$.

To implement σ^c we must combine top-12-fair deterministic assignments where $s_a, s_c \in \{4, 5\}$ and $s_b \in \{2, 3\}$, which leaves exactly three choices

	a	b	c
P_2	$\alpha\alpha\alpha\alpha$	$\alpha\beta$	$\beta\beta\beta\beta$
P_4	$\alpha\alpha\alpha\alpha$	$\alpha\beta\beta$	$\beta\beta\beta\beta$
P_5	$\alpha\alpha\alpha\alpha$	$\alpha\alpha\beta$	$\beta\beta\beta\beta$

Only P_2 is symmetric, it was already introduced in (5).

The lottery $L^c = \frac{1}{3}P_2 + \frac{1}{3}P_4 + \frac{1}{3}P_5$ (plus a uniform mixing of the α -s and of the β -s in their respective roles) is in this case the unique implementation of σ^c .

The expected total utility of the α -s is $\frac{1}{3}(7 + 6 + 10)$ so each agent's expects the utility 1.28. But the symmetric and top-12-fair assignment

$P_3 :$	a	b	c
	$\alpha\alpha\alpha\alpha$	$\alpha\alpha\beta\beta$	$\beta\beta\beta\beta$

collects more expected utility. After the uniform mixing inside each type, everyone gets 1.33.

The trade-off is between gathering more surplus at P_3 but allowing $s_b = 4$ so that the agents assigned to b are (ex ante and ex post) envious by a margin of 2 units of congestion, versus losing some surplus at the competitive lottery L^c while generating no ex ante envy and ex post envy only at the level of 1 unit of congestion.

Remark In the non congested assignment model another approach to randomisation extends each deterministic ordinal preference ordering to its (incomplete) stochastic dominance ordering of random allocations; More popular in this century; than cardinal vNM.

In the congested model this ordinal definition of competitiveness fails because it is vastly underspecified: if we fix an arbitrary congestion profile (s_x) we can force the PS algorithm ([7], [22], [4]) to implement it¹⁵ and easily select an envy free and ordinally efficient random assignment, but it is not clear what deeper principle should guide the choice of that congestion profile.

7 concluding comments

We submit that our competitive analysis applies to a rich family of congested assignment problems and delivers an efficient solution built upon compelling fairness principles.

main results Whether congestion is anonymous or weighted, the canonical guarantee eliminates for each participant all but the best $\frac{1}{m}$ -th quantile of the feasible allocations, where m is the number of posts (Propositions 1 and 2).

Competitiveness, when it exists in the deterministic version of either model, identifies an essentially single assignment in terms of congestion and welfare. Combining efficiency with natural ex ante and ex post fairness properties, it is then a compelling normative solution to the congested assignment problem (Propositions 3 and 4).

After eliciting vNM cardinal utilities rather than simpler ordinal preferences, the randomised competitive demand is unique (Lemmas 1 and 3) and implemented by a lottery over competitive deterministic assignments (Lemmas 2 and 4). In the

¹⁵For each post x we construct s_x copies of the allocation (x, s_x) and assign (randomly) these n objects to the n agents.

anonymous case the latter are close to each other, as well as approximately top- n -fair, competitive, and efficient (Theorem 1).

two open questions 1). The addition of post-specific lower and upper bounds on congestion is natural in many of the motivating examples discussed in the Introduction. Under anonymous congestion we can easily generalise the concept of top- n -fair guarantee but its interpretation is a bit more involved. For agent i the profile λ_i must satisfy (1) as well as the bounds: $\gamma_a^- \leq \lambda_{ia} \leq \gamma_a^+$ for all a . So the number of allocations that may or may not be accepted by λ_i is $\sum_a (\gamma_a^+ - \gamma_a^-)$ and λ_i captures the best $(n - \sum_a \gamma_a^-)$ of these. Then Proposition 1 goes through exactly as before.

Competitiveness is much harder to adapt. Its definition generalises when we have only upper bounds, but we lose the critical fixed point argument in Lemma 1 for the existence of a randomised competitive congestion profile.

With lower bounds it is not clear what the definition of competitiveness should be: how to account for the situation where a lower bound γ_a^- forces the agents to populate a post a that they unanimously loath?

2). A natural “dual” domain of preferences views congestion as strictly desirable: $(a, s + 1) \succ (a, s)$ for all a and s . The top-fairness idea still works (at least under anonymous congestion) but with much less bite. Interpret the report λ_i satisfying (1) as “agent i accepts (a, s) only if $s \geq \lambda_{ia}$ ”. These constraints over all agents are jointly feasible and allow each agent to veto at most $n - 1$ allocations, avoiding only the *lowest* $\frac{1}{m}$ quantile of their preferences. The proof mimicks that of Proposition 1 by switching a couple of signs.

The Definition 3 of a competitive assignment is unchanged but no longer identifies a unique solution for deterministic assignments (e. g. when everyone cares only about joining the most congested post). Proposing and justifying in the competitive spirit a fair randomised compromise in the “good congestion” model is a challenging question

References

- [1] Abdulkadiroglu A. and T. Sonmez. 2003. School Choice: A Mechanism Design Approach, *American Economic Review*, 93 (3), 729–747.

- [2] Banerjee S., Konishi, H. and Sonmez, T. 2001. Core in a Simple Coalition Formation Game, *Social Choice Welfare* 18, 135–153.
- [3] Biro P., Fleiner T., Irving R.W. and Manlove, D. F. 2010. The college admissions problem with lower and common quotas, *Theoretical Computer Science*, 411, 3136–3153.
- [4] Bogomolnaia A. 2015. Random assignment: redefining the serial rule, *Journal of Economic Theory*, 158, 308-318.
- [5] Bogomolnaia A. and M. O. Jackson. 2002. The stability of hedonic coalition structures, *Games & Economic Behavior*, 38,2, 201-230.
- [6] Bogomolnaia A., Le Breton M., Savvateev A. and S. Weber. 2007. Stability under unanimous consent, free mobility and core, *International Journal of Game Theory*, 35,2,185-204.
- [7] Bogomolnaia A. and H. Moulin. 2001. A new solution to the random assignment problem, *Journal of Economic Theory*, 100,2,201-230.
- [8] Boiteux M. 1960. Peak-Load pricing, *The Journal of Business*, 33, 2,157-179.
- [9] Budish, E. 2011. The combinatorial assignment problem: Approximate competitive equilibrium from equal incomes, *J. Political Economy*, 119, 6, 1061-1103.
- [10] Chen J., Guo G. and Y. Wen. 2023. Algorithms and Complexity for Congested Assignments, arXiv:/5288754 [cs:GT] 12 Dec 2023.
- [11] Cacchiani V., Iori M., Locatelli M. and Martello S. 2022. Knapsack problems, an overview of recent advances, Part II: Multiple, multidimensional, and quadratic knapsack problems, *Computers and Operations Research*, 143, 105693.
- [12] Caragiannis C., Kaklamanis C., Kanellopoulos P. and E. Papaioannou. 2008. Scheduling to maximize participation, *Theoretical Computer Science*, 402, 142-155.
- [13] Copland A. 2021. School choice and class size externalities. <http://dx.doi.org/10.2139/ssrn.4346860>

- [14] Correa J.R. and Stier-Moses N. E. 2010. Wardrop equilibria, *Wiley Encyclopedia of Operations Research and Management Science*, ed. James J. Cochran, John Wiley & Sons, Inc.
- [15] Ergin H. and Sönmez T. 2006. Games of school choice under the Boston mechanism, *Journal of Public Economics*, 90, 215-237.
- [16] Foley, D. K. 1966. *Resource allocation and the public sector*, Yale University.
- [17] Frank M. 1981. The Braess paradox, *Mathematical Programming*, 20, 283-302.
- [18] Green J. and J. J. Laffont. 1977. Characterization of Satisfactory Mechanisms for the Revelation of Preferences for Public Goods, *Econometrica*, 45, 2, 427-438.
- [19] Hafalir I. E., Yenmez M. B. and M. A. Yildirim. 2013. Effective affirmative action in school choice, *Theoretical Economics*, 8, 2, 325-363.
- [20] Hylland A. and R. Zeckhauser. 1979. The Efficient Allocation of Individuals to Positions, *Journal of Political Economy*, 87, 2, 293-311.
- [21] Kamada Y. and Kojima F. 2015. Efficient matching under distributional constraints: Theory and applications, *American Economic Review*, 105(1):67-99.
- [22] Katta A-K. and J. Sethuraman. 2006. A solution to the random assignment problem on the full preference domain, *Journal of Economic Theory*, 131, 1, 231-250.
- [23] Konishi H., Le Breton M., and S. Weber. 1997. Equilibrium in a Model with Partial Rivalry, *Journal of Economic Theory*, 72, 225-237.
- [24] Koutsoupias E. and C. H. Papadimitriou. 1999. Worst-case equilibria, *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science*, 404-413.
- [25] Milchtaich I. 1996. Congestion games with player-specific payoff functions, *Games & Economic Behavior*, 13, 1, 111-124.

- [26] Monderer D., and Shapley L. S. 1996. Potential Games, *Games & Economic Behavior*, 14, 1, 124-143.
- [27] Monte D. and Tumennasan, N. 2013. Matching with quorums, *Economics Letters*, 120(1):14–17
- [28] Moulin H. 2019. Fair Division in the Internet Age, *Annual Review of Economics*, 11, 407-41.
- [29] Papai S. 2000. Strategyproof assignment by hierarchical exchange, *Econometrica*, 68, 6, 1403-1433.
- [30] Phan W., Tierney R., and Y. Zhou. 2023. Crowding in School Choice, <https://wphan.wordpress.ncsu.edu>
- [31] Pigou A.C. 1920. *The Economics of Welfare*, London, Macmillan.
- [32] Ramsey, F. P. 1927. A Contribution to the Theory of Taxation, *Economic Journal*, 37, 47-61.
- [33] Rosenthal R. W. 1973. A Class of Games Possessing Pure-Strategy Nash Equilibria, *Int. J. Game Theory* 2, 65–67.
- [34] Roughgarden T. 2005. *Selfish Routing and the Price of Anarchy*, MIT Press, Cambridge Mass.
- [35] Roughgarden T. and E. Tardos. 2002. How bad is selfish routing? *Journal of the ACM*, 49(2), 236-259.
- [36] Shapley L. and H. Scarf. 1974. On Cores and Indivisibility, *Journal of Mathematical Economics*, 1, 23 37.
- [37] Steinhaus H. 1948. The Problem of Fair Division. *Econometrica*, 16(1), 101-4.
- [38] Varian H. 1974. Equity, Envy and Efficiency, *Journal of Economic Theory*, 29(2), 217-244.
- [39] Vickrey W. S. 1969. Congestion Theory and Transport Investment, *American Economic Review (Papers and Proceedings)*, 59, 251–260.

8 appendix

8.1 proof of Proposition 2

After proving that in any problem (A, N, \succeq, w) there exists at least one top- $\frac{1}{m}$ -fair assignment P , we state and prove the maximality property of the canonical guarantee.

Existence. As in Proposition 1 we combine a greedy algorithm and an induction on the number of agents. We can clearly get rid of the “inactive” posts s. t. $\lambda_{ia} = 0$ for all i , so we still write A for the set of active posts. We pick any post a and construct a set $S \subseteq A$ s. t.

$$\forall i \in S : \lambda_{ia} \geq w_S \text{ and } \forall j \in N \setminus S : \lambda_{ja} < w_S + w_j \quad (20)$$

Label the agents from 1 to n so that $\lambda_{ia} \geq \lambda_{(i+1)a}$ for all $i \in [n-1]$ where $\lambda_{1a} > 0$ because a is active. For any two disjoint subsets S, T in A we say that S rejects T if $\lambda_{ia} < w_S + w_T$ for some $i \in S$; otherwise we say that S accepts T . Note that if all labels in S are (weakly) smaller than all in T and S rejects T , then T rejects S as well; and S accepts T if T accepts S . We construct S recursively: in each step we either find S or add one agent to the provisional set.

step 1. If for all $j \geq 2$ the set $\{j\}$ rejects $\{1\}$ then $S = \{1\}$ meets (20) and we are done. Otherwise we pick the smallest label $\ell^1 \geq 2$ accepting $\{1\}$, which implies that $\{1\}$ accepts $\{\ell^1\}$ as well, and we form the provisional set $S^1 = \{1, \ell^1\}$. If $\ell^1 = n$ we are done by choosing S^1 so going into step 2 we have $\ell^1 < n$.

step $k+1$. The subset S has not yet been found therefore ℓ^k , the latest pick in S^k , is smaller than n . By construction $\lambda_{\ell^k a} \geq w_{S^{k-1}} + w_{\ell^k} = w_{S^k}$ so $\lambda_{ia} \geq w_{S^k}$ for all $i \in S^k$. Moreover all agents $j < \ell^k$ outside S^k have rejected some earlier $S^{k'}$, so they also reject the larger set S^k .

If all agents $j > \ell^k$ reject S^k as well we are done by choosing S^k . Otherwise we pick the smallest label ℓ^{k+1} after ℓ^k s.t. $\{\ell^{k+1}\}$ accepts S^k : this implies that S^k accepts $\{\ell^{k+1}\}$ as well (recall $\lambda_{\ell^k a} \geq \lambda_{\ell^{k+1} a}$) so we set $S^{k+1} = S^k \cup \{\ell^{k+1}\}$ and we have $\lambda_{ia} \geq w_{S^{k+1}}$ for all $i \in S^{k+1}$. We are done if $\ell^{k+1} = n$ otherwise we go to the next step. When this process stops we have found \tilde{S} meeting (20).

We assign \tilde{S} to post a , and consider the residual problem in $\tilde{A} = A \setminus \{a\}$, $\tilde{N} = N \setminus \tilde{S}$ with total congestion $\tilde{W} = W - w_S$. For each $j \in \tilde{A}$ such that $\lambda_{ja} > 0$

inequality (20) and equation (6) imply

$$\lambda_{j\tilde{A}} > \lambda_{jA} - (w_S + w_j) = W - w_S + ([[\lambda_j]] - 2)w_j = \widetilde{W} + ([[\tilde{\lambda}_j]] - 1)w_j$$

where in $\tilde{\lambda}_j$ we drop λ_{ja} . For each $j \in \tilde{A}$ such that $\lambda_{ja} = 0$ inequality (??) has no bite and equation (6) gives

$$\lambda_{j\tilde{A}} = W + ([[\lambda_j]] - 1)w_j > \widetilde{W} + ([[\lambda_j]] - 1)w_j = \widetilde{W} + ([[\tilde{\lambda}_j]] - 1)w_j$$

The induction argument shows that in the reduced problem there is a top- $\frac{1}{m-1}$ -fair assignment of $N \setminus S$ to \tilde{A} .

The **Maximality** property: *Fix A, W , an agent i^* with weight w_{i^*} and an arbitrary $(W - w_{i^*})$ -prefix λ_{i^*} ((6)). Then for any a s. t. $\lambda_{i^*a} > w_{i^*}$ there is a set of agents M not containing i^* with weights w_i s. t. $\sum_{i \in M} w_i = W - w_{i^*}$ and a prefix λ_i ((6)) for each $i \in M$, s. t. in any top- n -fair assignment of the $M \cup \{i^*\}$ problem agent i^* is at a post a where the congestion is arbitrarily close to λ_{i^*a} .*

Proof: We fix W, i^* , the $(W - w_{i^*})$ -prefix λ_{i^*} and a as in the statement. The set M contains one agent i_b for each post b s. t. $\lambda_{i^*b} \geq w_{i^*}$, in particular one i_a . The size of M is $[[\lambda_{i^*}]]$. The weights are $w_{i_b} = \lambda_{i^*b} - w_{i^*} + \varepsilon$ if $b \in M \setminus a$ and $w_{i_a} = \lambda_{i^*a} - w_{i^*} - ([[\lambda_{i^*}]] - 1)\varepsilon$, where $\varepsilon > 0$ is small enough that $w_{i_a} > 0$. Clearly $w_{i^*} + \sum_{b \in M} w_{i_b} = W$. Suppose now that each i_b (including i_a) is single-minded on post b : $\lambda_{i_b b} = W, \lambda_{i_b x} = 0$ for $x \neq b$. Consider a top- $\frac{1}{m}$ -fair assignment of $M \cup \{i^*\}$ to A : for each i_b other than i_a the inequality $w_{i_b} + w_{i^*} > \lambda_{i^*b}$ implies that i_b is alone at b , therefore i^* share a with i_a where the congestion $\lambda_{i^*a} - ([[\lambda_{i^*}]] - 1)\varepsilon$ is arbitrarily close to λ_{i^*a} . ■

8.2 proof of Proposition 4

It mimicks that of Proposition 3. Given the assignment $P = (S_x)_{x \in A}$ for simplicity we write the congestion w_{S_x} simply as s_x . Now agent i 's competitive demand at the congestion profile s is

$D(i; s) = \{a | (a, s_a \vee w_i) \succeq_i (x, s_x \vee w_i) \text{ for all } x \in A\}$, and P is competitive iff $a \in D(i; s)$ whenever $i \in S_a$.

Statement i) Unique competitive congestion. At first we do not assume in (A, N, \succeq, w) either semi-strict preferences or a crowded competitive assignment.

Fix $P = (S_x)_{x \in A}$ and $P^* = (S_x^*)_{x \in A}$ both competitive and such that $s \neq s^*$. Partition A as $B^* \cup C \cup B$ where $B^* = \{a : s_a < s_a^*\}$, $B = \{a : s_a^* < s_a\}$, $C = \{a : s_a^* = s_a\}$, and B, B^* are both non empty. We define in two equivalent ways the set $A^* = \{a \in B^* : s_a^* > w_i \text{ for all } i \in S_a^*\} = \{a \in B^* : |S_a^*| \geq 2\}$.

We claim first (as in the proof of Proposition 3) that if the demand $D(i; s^*)$ of some agent i intersects A^* , then $D(i; s) \subseteq A^*$. Suppose, to the contrary $a \in D(i; s^*) \cap A^*$ while $b \in D(i; s) \cap A \setminus A^*$. By definition of the demands and of B^* we have

$$(a, s_a^*) \succeq_i (b, s_b^* \vee w_i) \text{ and } (b, s_b \vee w_i) \succeq_i (a, s_a \vee w_i) \succ_i (a, s_a^*)$$

where the last relation is strict because $s_a \vee w_i < s_a^*$. Note that we cannot replace $s_b \vee w_i$ by s_b because i may not be in S_b .

If $b \notin B^*$ we have $s_b \vee w_i \geq s_b^*$ so we add $(b, s_b^* \vee w_i) \succeq_i (b, s_b \vee w_i)$ to the above preferences to get a contradiction. If $b \in B^*$ then $|S_b^*| = 1$ because $b \notin A^*$, therefore $s_b < s_b^* = w_i$ so that $(b, s_b \vee w_i) = (b, s_b^* \vee w_i)$ and we reach again a contradiction. The claim is proved. As in the proof of Proposition 1 we check next that A^* must be empty: every $i \in \cup_{a \in A^*} S_a^*$ is s. t. $D(i; s) \subseteq A^*$ therefore $\sum_{a \in A^*} s_a \geq \sum_{a \in A^*} s_a^*$, which contradicts $s_a^* > s_a$ in B^* . So for every $a \in B^*$ we have $S_a^* = \{i\}$. Exchanging the roles of P and P^* we see that $S_b = \{i\}$ for all $b \in B$.

If at least one competitive assignment is crowded we take it as one of P, P^* : this contradicts the existence of B and B^* and we are done. We continue the proof when we only know that preferences are semi-strict. Writing $T^* = \cup_{a \in B^*} S_a^*$ and $T = \cup_{a \in B^*} S_a$ we have $w_{T^*} = \sum_{a \in B^*} s_a^* > \sum_{a \in B^*} s_a = w_T$, implying that $T^* \setminus T$ is not empty: it contains some agent $i \in S_a^* \cap S_b$ where $a \in B^*$ and $b \in C \cup B$. By definition of the demands and of the partition of A we have

$$(a, w_i) = (a, s_a^*) \succeq_i (b, s_b^* \vee w_i) \text{ and } (b, s_b) \succeq_i (a, s_a \vee w_i) = (a, w_i)$$

where the last equality is from $s_a < s_a^* = w_i$. If $b \in B$ we have $w_i = s_b > s_b^*$ therefore $(b, s_b^* \vee w_i) \simeq_i (b, s_b)$; this indifference still holds if $b \in C$ because in that case $s_b^* = s_b \geq w_i$: we conclude that all preferences above are indifferences in particular $(a, w_i) \simeq_i (b, s_b)$ which contradicts the semi-strictness of preferences.

Statement i) Unique competitive welfare. Fix $P = (S_x)_{x \in A}$ and $P^* = (S_x^*)_{x \in A}$ both competitive; we just proved $s = s^*$. An agent i is in some $S_a^* \cap S_b$ (where

$a = b$ is possible); by definition of competitiveness and the inequalities $w_i \leq s_a, s_b$ we have $(a, s_a) \succeq_i (b, s_b)$ and $(b, s_b) \succeq_i (a, s_a)$.

Statement ii) Assume P is competitive and Pareto dominated by Q . Pick any agent i who is in S_a at P and in T_b at Q ($a = b$ is possible). We have

$$(b, t_b) \succeq_i (a, s_a) \succeq_i (b, t_b \vee w_i) \quad (21)$$

therefore $t_b \vee w_i \geq t_b$. Suppose $s_b < t_b$: taking into account $w_i \leq t_b$, the inequality $s_b \vee w_i \geq t_b$ implies $w_i = t_b > s_b$, and the two preferences in (21) are indifferences, in particular $(b, w_i) \simeq_i (a, s_a)$ a contradiction of semi-strict preferences if $a \neq b$. If $a = b$ then (21) gives $(b, t_b) \succeq_i (b, s_b)$ contradicting $s_b < t_b$.

We have shown $s_b \geq t_b$ everywhere on the support of Q , so these are all equalities. By replacing s_b by t_b in (21) we get $(b, t_b) \simeq_i (a, s_a)$ a final contradiction.

Statement iii) Fix a competitive assignment P , a post a in its support and an agent $i \in S_a$. We use the preferences $(a, s_a) \succeq_i (x, w_{S_x} \vee w_i)$ for all $x \neq a$ to evaluate the size of the set of allocations agent i may prefer to (a, s_a) : at post a the size is $s_a - w_i$; at x such that $w_{S_x} \geq w_i$ the size is $w_{S_x} - w_i$; and 0 at x such that $w_{S_x} \leq w_i$. So the total is at most $W - w_i$. ■

8.3 fractional competitiveness under weighted congestion

Switching to cardinal vNM utility functions we only partially emulate the results of the previous section. Two of the key results are preserved: the existence of a competitive fractional congestion profile and its qualified uniqueness (Lemma 3 below); the implementation of the competitive congestion by one or more lotteries over deterministic assignments (Lemma 4 below). However we lose the rounding and approximation results of Lemma 2 and Theorem 1.

The congestion profile $\sigma = (\sigma_a)_{a \in A}$ is in the simplex $\Delta^{\mathbb{R}}(A, W)$ and each agent i has a cardinal vNM utility function $u_i(a, z)$ over $A \times [w_i, W]$, continuous and strictly increasing in z . As before if $\sigma_a < w_i$ the congestion price of post a to agent i is in fact w_i so we do not define $u_i(a, z)$ in the interval $0 \leq z < w_i$. Given a problem (A, N, u, w) agent i 's competitive demand at σ is

$D(u_i, \sigma) = \arg \max_{x \in A} u_i(x, \sigma_x \vee w_i)$ and the fractional congestion profile $\sigma \in \Delta^{\mathbb{R}}(A; W)$ is competitive (Comp) iff $\sigma \in \sum_{i \in N} w_i \cdot \Delta^{\mathbb{R}}[D(u_i, \sigma); 1]$.

The (not necessarily unique) corresponding semi-stochastic matrix Π meets (18), but (17) is replaced by

$$\sigma_a = \sum_{i \in N} w_i \pi_{ia} \text{ for all } a \in A \quad (22)$$

The same fixed point argument as in Lemma 1 proves that a solution of (22) exists. To prove that it is unique we must strengthen the crowdedness property of Definition 5 and prove a weaker uniqueness statement than in Proposition 4 (subsection 5.2).

Definition 8 *The fractional congestion profile σ is f-crowded in problem (A, N, u, w) iff for all a in its support ($\sigma_a > 0$) and each agent i demanding a ($a \in D(u_i, \sigma)$) we have $\sigma_a > w_i$.*

Lemma 3: *In any problem (A, N, \succeq, w) there is at most one f-crowded competitive congestion profile.*

Proof Let σ, σ^* be two different f-crowded competitive congestion profiles, and B^* the non empty set of posts s. t. $\sigma_a^* > \sigma_a$. We claim that if $D(u_i, \sigma^*)$ contains a post $a \in B^*$ for some agent i then $D(u_i, \sigma) \subseteq B^*$. If, to the contrary, $D(u_i, \sigma)$ contains b outside B^* we get the familiar inequalities $u_i(a, \sigma_a^*) \geq u_i(b, \sigma_b^* \vee w_i)$ and $u_i(b, \sigma_b) \geq u_i(a, \sigma_a \vee w_i) > u_i(a, \sigma_a^*)$ where the strict inequality is from $\sigma_a < \sigma_a^*$ and the f-crowding assumption. If $b \notin B^*$ we have $\sigma_b^* \vee w_i \leq \sigma_b$ and derive a contradiction from $u_i(b, \sigma_b^* \vee w_i) \geq u_i(b, \sigma_b)$.

Pick now some semi-stochastic matrices Π and Π^* meeting (22) for σ and σ^* respectively, and use the property $D(u_i, \sigma^*) \cap B^* \neq \emptyset \implies D(u_i, \sigma^*) \subseteq B^*$ to compute for any $i \in N$: $\sum_{a \in A^*} \pi_{ia}^* > 0 \implies \sum_{a \in A^*} \pi_{ia}^* \leq \sum_{a \in A^*} \pi_{ia}$, which implies $\sum_{a \in A^*} \sigma_a^* = \sum_{i \in N} w_i (\sum_{a \in A^*} \pi_{ia}^*) \leq \sum_{i \in N} w_i (\sum_{a \in A^*} \pi_{ia}) = \sum_{a \in A^*} \sigma_a$, the final contradiction. ■

The weighted congestion analog of Lemma 2 (section 6 for the anonymous congestion case) is next.

Lemma 4 *In any problem (A, N, u, w) we can implement any competitive congestion profile σ^c by one or more list $(\{P^k\}_{k=1}^K, \mathcal{L})$ of K deterministic assignments P^k together with a lottery \mathcal{L} over $[K]$, such that their expected congestion is σ^c and they always assign the agents in their competitive demands:*

$$i \in S_a^k \implies a \in D(u_i, \sigma^c).$$

Proof Pick Π realising the congestion σ^c in (22) and apply the following version of the Birkhoff theorem: every semi-stochastic matrix Π is the convex combination of deterministic matrices (all entries are 0 or 1) of which the support is contained in that of Π . ■

A final example illustrates the limits of this result.

Example 7: *Two posts, three agents* $m = 2, n = 3$ and $W = 21$. The weights, $(W - w_i)$ -prefixes, and the two top- $\frac{1}{2}$ -fair assignments are

	a	b	w		a	b
α	16	15	10		P_1	$\alpha \quad \beta\gamma$
β	16	15	10		P_2	$\beta \quad \alpha\gamma$
γ	9	13	1			

If the preferences are represented by the slack utilities (7) neither P_1 or P_2 is competitive because both α and β prefer $(a, 10)$ to $(b, 11)$.

The fractional competitive congestion is $\sigma^c = (11, 10)$ because $D(\alpha, \sigma^c) = D(\beta, \sigma^c) = \{a, b\}$. It can be implemented by two lotteries $\mathcal{L}_1, \mathcal{L}_2$ combining P_1, P_2 with three not top- $\frac{1}{2}$ -fair assignments:

	a	b
P_3	$\alpha\gamma$	β
P_4	$\beta\gamma$	α
P_5	$\alpha\beta$	γ

Specifically: $\mathcal{L}_1 = \frac{1}{4}P_1 + \frac{1}{4}P_2 + \frac{1}{4}P_3 + \frac{1}{4}P_4$, and $\mathcal{L}_2 = \frac{9}{20}P_1 + \frac{9}{20}P_2 + \frac{1}{10}P_5$. The former keeps congestion close to top- $\frac{1}{2}$ -fair, but violates this property half of the time; the latter only violates it with probability $\frac{1}{10}$ but then much more severely.