

האוניברסיטה העברית בירושלים
THE HEBREW UNIVERSITY OF JERUSALEM

**NONCOOPERATIVE MARKET ALLOCATION
AND THE FORMATION OF DOWNTOWN**

By

**YANNAI A. GONCZAROWSKI,
and MOSHE TENNENHOLTZ**

Discussion Paper # 663 (March 2014)

מרכז לחקר הרציונליות

**CENTER FOR THE STUDY
OF RATIONALITY**

Feldman Building, Givat-Ram, 91904 Jerusalem, Israel
PHONE: [972]-2-6584135 FAX: [972]-2-6513681
E-MAIL: ratio@math.huji.ac.il
URL: <http://www.ratio.huji.ac.il/>

Noncooperative Market Allocation and the Formation of Downtown

YANNAI A. GONCZAROWSKI, Hebrew University of Jerusalem and Microsoft Research
MOSHE TENNENHOLTZ, Microsoft Research and Technion — Israel Institute of Technology

Can noncooperative behaviour of merchants lead to a market allocation that *prima facie* seems anticompetitive? We introduce a model in which service providers aim at optimizing the number of customers who use their services, while customers aim at choosing service providers with minimal customer load. Each service provider chooses between a variety of levels of service, and as long as it does not lose customers, aims at minimizing its level of service; the minimum level of service required to satisfy a customer varies across customers. We consider a two-stage competition, in the first stage of which the service providers select their levels of service, and in the second stage — customers choose between the service providers. (We show via a novel construction that for any choice of strategies for the service providers, a unique distribution of the customers' mass between them emerges from all Nash equilibria among the customers, showing the incentives of service providers in the two-stage game to be well defined.) In the two-stage game, we show that the competition among the service providers possesses a unique Nash equilibrium, which is moreover super strong; we also show that all sequential better-response dynamics of service providers reach this equilibrium, with best-response dynamics doing so surprisingly fast. If service providers choose their levels of service according to this equilibrium, then the unique Nash equilibrium among customers in the second phase is essentially an allocation (i.e. split) of the market between the service providers, based on the customers' minimum acceptable quality of service; moreover, each service provider's chosen level of service is the lowest acceptable by the entirety of the market share allocated to it. Our results show that this seemingly-cooperative allocation of the market arises as the unique and highly-robust outcome of noncooperative (i.e. free from any form of collusion), even myopic, service-provider behaviour. The results of this paper are applicable to a variety of scenarios, such as the competition among ISPs, and shed a surprising light on aspects of location theory, such as the formation and structure of a city's central business district.

Key Words and Phrases: Game Theory, Congestion Games, Location Theory, Two-Stage Competition

1. INTRODUCTION

1.1. Setting

1.1.1. Shoppers. Consider the downtown area of the fictional city of Metropolis, the wine capital of the world. At its heart lies Metropolis Central Station. Every morning, shoppers from throughout the Metropolis metropolitan area (and beyond) disembark the train at Metropolis Central, at the vicinity of which many wine shops are located, and go about their wine-shopping errands. Each shopper is interested in purchasing a single bottle of wine, and is willing to walk at most d minutes (a shopper-dependant real value) in each direction in order to get it. Assume for the time being that the qualities of the various wines available on the market are indistinguishable (or alternatively, that our shoppers do not care about the quality of the wine that they get); we remove this assumption in Section 1.1.4 below. Therefore, each shopper's sole consideration is that the bottle of wine that she buys be as exclusive as possible, i.e. she prefers to get her wine at the shop that sells the fewest bottles of wine throughout the day (so that it can be considered a "boutique wine"), as long as it is no more than d minutes away from Metropolis Central, of course. A *Nash equilibrium* among the shoppers is therefore an assignment of the shoppers to shops, s.t. for each shopper with walking-time limit d , no shop exists within d walking minutes of Metropolis Central that sells fewer bottles of wine throughout the day than the shop assigned to this shopper.

We consider a scenario with finitely-many shops and continuously-many shoppers,

Authors' addresses: Y. A. Gonczarowski, Einstein Institute of Mathematics, Rachel and Selim Benin School of Computer Science and Engineering and Center for the Study of Rationality, Hebrew University of Jerusalem, Israel; and Microsoft Research, *Email*: yannai@gonch.name; M. Tennenholtz, Microsoft Research and Technion — Israel Institute of Technology, *Email*: moshet@microsoft.com.

the distribution of d among whom is given by an arbitrary finite Lebesgue measure. Preparing the ground for the main results of this paper, which follow below, in Section 4 we use a novel construction to show the following result (similar in spirit to other results regarding congestion games and crowding games).

THEOREM 1.1 (INFORMAL VERSION OF THEOREMS 4.5, 4.10 AND 4.11, COROLLARY 4.12, AND ALGORITHM 1).

- (1) *A Nash equilibrium among the shoppers exists. Furthermore, there exists such a Nash equilibrium for which the strategies can be computed in $O(n^2)$ time, where n is the number of shops.*
- (2) *The number of bottles of wine sold at each shop is the same across all Nash equilibria, and for each shopper, the exclusivity of the wine that she buys (i.e. the number of bottles sold throughout the day at the store selling this wine) is the same across all Nash equilibria. All these amounts can be computed together in $O(n^2)$ time.*

1.1.2. Wine sellers. Obviously, each wine seller would like to locate her store in a way that would maximize its sales volume. (By Theorem 1.1, this sales volume is well defined, assuming a Nash equilibrium among the shoppers.) That being said, as real-estate prices rise the closer (in walking time) a shop is to Metropolis Central, each wine seller would like to place her store the farthest possible from the station, as long as this does not hurt sales. As we think of sales as indicative of daily income, and of real-estate cost as a one-time expense, we have that each wine seller would like to place her shop in a location that first and foremost maximizes sales, and only then (as a tie-breaking rule among locations that generate the same volume of sales) the farthest away (in walking time) from Metropolis Central. In Section 5.2, we show the following — the first of our main results regarding this two-stage competition.

THEOREM 1.2 (INFORMAL VERSION OF THEOREM 5.28 AND PROPOSITION 5.30).

- (1) *The following is a super-strong¹ pure-strategy Nash equilibrium among n wine sellers: the shop that is closest to Metropolis Central is located in the farthest location from the station that is accessible by all shoppers, the second-closest shop is located at the first n -tile of walking distances (i.e. the farthest place accessible by all but $1/n$ of shoppers), the third-closest shop is located at the second n -tile of walking distances (i.e. the farthest place accessible by all but $2/n$ of shoppers), etc.*
- (2) *All shops have identical sales volumes in this Nash equilibrium.*
- (3) *No other (not necessarily super strong) Nash equilibria (up to permutations on the shops) exist, not even in mixed strategies.*

We show the robustness of the Nash equilibrium defined in Theorem 1.2 even further by considering dynamics among wine sellers. A *sequential* best-response dynamic is a process starting with arbitrary shop locations, and in which at each turn an arbitrary wine seller relocates her shop to a location that, *ceteris paribus*, maximizes her preferences (we show that such a location, or rather such a distance, always exists for every possible measure on walking-time limits for shoppers); we assume that each wine seller is allowed to relocate her shop infinitely often. A *round* in a best-response dynamic is a sequence of consecutive steps in which each wine seller is allowed to relocate her shop at least once. Finally, a *sequential δ -better-response* dynamic is a sequential dynamic in which each relocation need not necessarily maximize the wine seller's preferences, as long as it increases her sales volume by at least δ of the entire market share (since we consider continuously-many shoppers, we demand a δ -improvement in

¹See Theorem 5.6 in Section 5.1 and the preceding discussion for the definition of super-strong equilibrium.

order to avoid improvements *à la* Zeno’s “Race Course” paradox.) In Section 5.2, we show the following main result.

THEOREM 1.3 (INFORMAL VERSION OF COROLLARIES 5.37 AND 5.38).

- (1) *For every $\delta > 0$, every sequential δ -better-response dynamic reaches a Nash equilibrium in finitely-many steps, and remains constant from that point onward.*
- (2) *Every best-response dynamic reaches a Nash equilibrium in at most n rounds.*

We also analyse dynamics in which more than one wine seller relocate their shops simultaneously. (See Corollary 5.35 and Theorems 5.15 and 5.22.)

1.1.3. Market Allocation. In the equilibrium defined in Theorem 1.2, the $1/n$ of shoppers with smallest walking-time limits shop at the store closest to Metropolis Central, whose chosen proximity to Metropolis Central is the farthest that still accommodates this entire $1/n$; the next $1/n$ of shoppers shop at the store second-closest to Metropolis Central, whose chosen proximity to Metropolis Central is the farthest that still accommodates this entire $1/n$; and so forth. Essentially, **the market is allocated (i.e. split) among the various wine sellers based on the willingness of a shopper to venture far from Metropolis Central, and each wine seller chooses the farthest location accessible by the entirety of its allocated market share.** Theorem 1.3 shows that this seemingly-cooperative market allocation among the various wine sellers arises as the unique possible outcome, not as a result of anticompetitive practices, but rather as a result of noncooperative dynamics, each wine seller only looking to myopically maximize her preferences at every step; no signalling (via e.g. location choice) or any other collusive or cooperative “trick” whatsoever is used in order to reach and maintain this market allocation. We further demonstrate this in Section 1.2.

1.1.4. Heterogeneous Wines. Recall the above assumption that the qualities of the various wine types are indistinguishable. Consider that this is no longer the case, i.e. that some wines may be known to be of superior types, are more extravagantly packaged, or have some other attractive quality. In such a case, shoppers may be willing to compromise on “exclusivity” in favour of superior quality. More generally, instead of purchasing wine of minimum circulation, each shopper would like to minimize a wine-seller-dependent increasing function of this circulation, e.g. shoppers may wish to maximize the quotient of quality and circulation. We discuss such scenarios in Section 6. All of our results, and in particular Theorems 1.1 to 1.3, readily generalize to these scenarios as well, with only quantitative rather than qualitative changes; e.g. the unique market-share allocation in a merchant-equilibrium (which we precisely characterize) is no longer necessarily the allocation of $1/n$ of the shoppers to each shop. In particular, we still obtain that the unique possible noncooperative outcome is for the market to be allocated (split) among the various wine sellers based on the shoppers’ types, i.e. each shopper shopping at the store closest to Metropolis Central has a smaller walking-time limit than any of those shopping at the store second-closest to Metropolis Central, each of whom in turn having a smaller walking-time limit than all of those shopping at the third-closest store, etc., and each wine store chooses the farthest location accessible by the entirety of its allocated market share.

1.2. Alternative Interpretations

It is worthwhile to point out that our framework captures far more than merely the store-location scenario introduced above, by thinking of the walking time from Metropolis Central Station to a wine store as a quality of service (QoS) of sorts of this store — the closer a shop is to Central Station, the higher the quality of service that it provides. We now give a few brief examples of other possible applications stemming

from this insight; in each example, QoS is given a different meaning, which, in turn, results in a different meaning of market allocation (split) based on acceptable QoS. These examples provide insight into the breadth of meanings that can be captured by the idea of QoS, and into the meaning of market allocation (split) based on acceptable QoS. We henceforth use the more generic term *producers* to refer to e.g. wine sellers, and *consumers* to refer to e.g. shoppers. We start with a main applicative example.

Example 1.4 (Home Internet / Cellular Market; QoS=Low Latency). Consider a scenario in which the producers are internet service providers (ISPs), and the consumers are customers on the market for home internet. (Alternatively, one could think of producers as cellular operators, and of consumers as customers on the market for cellular service.) An ISP is typically characterized by two parameters: the *de facto* bandwidth (i.e. speed) of the connection that it provides, and the latency of its infrastructure. (Assume that the monthly price for a home internet connection is the same for all ISPs, and that this price is precisely what each customer can afford to pay for an internet connection.) Each customer is looking for one internet connection, and finds acceptable any connection with a latency of at most d milliseconds (a customer-dependent value — e.g. real-time gaming applications are known to be quite latency-sensitive); a customer would like to get the fastest connection possible, given that her latency constraint is met. If each ISP has the same total (i.e. overall) bandwidth, then the speed of the connection of a single customer subscribed to an ISP is inversely proportional to this ISP's number of subscribers, and so obtaining the fastest connection possible is equivalent to subscribing to a least-subscribed-to ISP. (Generalizing, we may imagine that some ISPs may have different total bandwidths than others, while some other ISPs may purchase additional total bandwidth as their subscriber pool grows; in such a scenario, in order to surf with greatest speed, each customer would prefer to subscribe not necessarily to a least-subscribed-to ISP, but rather to an ISP from whom the customer would receive the fastest connection, a calculation depending both on the total-bandwidth characteristics of the specific ISP and on its number of subscribers. As discussed above in the context of heterogeneous wines, we handle such scenarios as well.) Finally, each ISP would like to maximize its number of subscribers, but would like to erect the cheapest-possible (i.e. with highest latency) infrastructure, as long as this does not reduce its market share.²

In this scenario, the unique possible noncooperative outcome is the allocation of each $1/n$ (this fraction varies here and henceforth if the total-bandwidth characteristics of ISPs are heterogeneous) of customers (based on acceptable latency levels) to a different ISP: the $1/n$ of customers with strictest latency demands is allocated to one ISP (whose network is constructed to have the highest latency acceptable by this entire $1/n$), the second $1/n$ is allocated to a second ISP (whose network is in turn constructed to have the highest latency acceptable by this entire $1/n$), and so forth.

Example 1.5 (Postgraduate Scholarships; QoS=Earlier Award Time). In Israel, there are three top-tier scholarships for science postgraduate students, all offering similar monetary support and all alike in prestige. These scholarships are mutually exclusive, in the sense that not only can no student be awarded more than one of them, but furthermore, it is not allowed to apply for more than one of them at any given year (anyone doing so is immediately disqualified from all three). Consider a rent-paying

²Assume that the investment in infrastructure is a one-time expense, and that the infrastructure upkeep cost is independent of the latency. Alternatively, considering cellular operators in lieu of ISPs, one may consider a situation in which at the end of the current licensing period, the license will be renewed only to the most-subscribed-to cellular operators, creating an incentive to first and foremost maximize the number of subscribers, even at the expense of higher infrastructure cost and upkeep.

student (a consumer) faced with the decision of which scholarship (producer) to apply to; if this student has no prior regarding neither her quality nor that of other applicants, it is in her best interest to subscribe to the least-sought-after scholarship, as long as it is awarded no later than when this student's yearly rent payment is due. (Once again, one may consider scholarships of heterogeneous prestige or monetary prizes, in which case students have a tradeoff between scholarship quality and number of competitors.) Each scholarship fund would like to maximize the quality of its fellows, which, having no prior regarding student qualities either, essentially means maximizing its number of applicants; as long as this quality (equivalently, the number of applicants) is not compromised, a scholarship fund would like to award the scholarship as late as possible throughout the year, in order to earn more interest at the bank. (Assume that bank interests are progressive, and so negligible for any single student, but substantial for a fund.)

In this scenario, the unique possible noncooperative outcome is the allocation of each third (when all scholarships are as attractive) of students (based on rent payment dates) to a different scholarship: the third with earliest rent payment dates are allocated to one scholarship (whose award date is chosen to be the earliest of the rent payment dates of its applicants), the second third is allocated to a different scholarship (whose award date is in turn chosen to be the earliest of the rent payment dates of its applicants), and the third with latest rent payment dates are allocated to the remaining scholarship (whose award date is chosen to be the earliest of the rent payment dates of its applicants).

1.3. The Formation of Main Street

Now that we have presented the main results of this paper, let us briefly return to Downtown Metropolis. While we have characterized the distance of each wine shop from Metropolis Central, the direction from Metropolis Central to each such shop can be arbitrary. Not for long, though. Merchants from the nearby town of Smallville, the extra-extra-extra-virgin-olive-oil capital of the world, wishing to widen the visibility of their product, have started moving their stores to Downtown Metropolis as well. Now that Metropolis has become both the wine- and the extra-extra-extra-virgin-olive-oil capital of the world, each shopper arriving at Metropolis Central would like to purchase not only a bottle of wine, but also a bottle of olive oil. Nonetheless, the walking-time limit of each shopper does not change — each shopper is still willing to walk at most $2d$ minutes in order to obtain both products. (This indeed introduces no change, as each shopper was previously willing to walk d minutes *in each direction*.) As with wine, each costumer prefers to minimize a seller-dependent function of the circulation of the type of olive oil that she purchases, as long as her walking-time constraint is met. (One may again consider e.g. the case in which one would like to maximize the quotient of quality to circulation, optimizing some form of tradeoff between quality and “boutiqueness”.) Olive-oil merchants have preferences similar to those of wine sellers.

In Section 7, we show that under these conditions, generally the unique stablest outcome is for all shops to be placed **on the same ray** originating at Metropolis Central (with the distance of each store from Metropolis Central set as before, as if its good type were the only one on the market). This should be compared with the structure of many old European towns, at the heart of which lies the old stone-cobbled main street, on one end of which (as opposed to at the middle of which) lies the main town church.

2. RELATED WORK

Our consumer games are a form of congestion games. Congestion games with finitely-many players have been introduced by Rosenthal [1973]; in fact, the term has been coined in a paper by Monderer and Shapley [1996], titled “Potential Games”, where it

is shown that a game has a potential iff it is a congestion game. While the discussion there refers to atomic games with finitely-many players, work in computer science and game theory also deals with nonatomic games, in which there may be a continuum of players (see e.g. [Roughgarden and Tardos 2002] for work in CS that uses such a model). Holzman and Law-Yone [1997] and Holzman [2003] look at restrictions on strategy sets of congestion games; one way to view our consumer games is as a special form of restricted congestion games defined for general measure functions on agents' types, capturing their possible strategy sets. As it turns out, this set of games possess many desired game-theoretic properties. The actual games that we study are in fact two-stage games, where the second stage is a congestion game as discussed above; the first stage can be viewed as a form of facility-location game, where the main aim of producers is to select a location to be selected by as many consumers as possible. This resembles the literature on location theory initiated by Hotelling [1929], although the utility function of the service providers in our setting is different, and allows for fine preferences based on distance from a location most preferred by consumers. Given the above, our model can be viewed as a novel combination of facility-location games among producers with congestion games among consumers.

Another type of related literature deals with scheduling and queuing with multiple machines, where the jobs choose among available services and the level of service they receive depends on the selections by other jobs. Recently, two-stage games in these contexts have been studied, consisting of a strategic selection by machines between queuing policies [Ashlagi et al. 2013] or scheduling policies [Ashlagi et al. 2010], followed by a strategic selection by jobs between the various selected policies. Our work introduces a novel type of a two-stage scenario, which may be considered as somewhat related. More remotely is the literature on competing mechanisms in the context of auctions, which employs such two phase setting, but in a very different context of mechanism design with money (see e.g. [McAfee 1993]).

The proof of Theorem 1.1 draws its intuition from an analogy to a hydraulic system of communicating vessels (see Fig. 1). Kaminsky [2000] uses an analogy to quite a different system of communicating vessels to solve rationing problems; his motivation is quite different, and involves extending bilateral rationing rules. While Kaminsky uses a set of two-way communicating vessels, we use a set of one-way communicating vessels. In this context, the problem of finding a Nash equilibrium among consumers may be regarded as a rationing problem with certain "reserves" for producers with high quality of service. Our treatment, especially in light of the discussion in Section 6 (see in particular Fig. 2), also sheds new light on rationing problems, as congestion games of sorts among a continuum of good-fragments.

3. NOTATION

Definition 3.1 (Notation).

— (*Simplex*). For a finite set S and a nonempty subset $S' \subseteq S$, we define

$$\Delta^{S'} = \left\{ s \in [0, 1]^S \mid \sum_{j \in S'} s_j = 1 \ \& \ \forall j \in S \setminus S' : s_j = 0 \right\}.$$

(The set S will be clear from context.)

— (*Naturals*). We denote the natural numbers by $\mathbb{N} \triangleq \{0, 1, 2, \dots\}$.

— (*Nonnegative Reals*). We denote the nonnegative reals by $\mathbb{R}_{\geq} \triangleq \{r \in \mathbb{R} \mid r \geq 0\}$.

— (*Maximizing Arguments*). Given a set S and a function $f : S \rightarrow \mathbb{R}$ that attains a maximum value on S , we denote the *set* of arguments in S maximizing f by $\arg \text{Max}_{s \in S} f(s) \triangleq \{s \in S \mid f(s) = m\}$, where $m \triangleq \text{Max}_{s \in S} f(s)$.

- For every $n \in \mathbb{N}$, we define $\mathbb{P}_n \triangleq \{0, 1, \dots, n-1\}$.
- Given a tuple $\bar{t} = (t_0, \dots, t_{n-1}) \in S^{\mathbb{P}_n}$ for some set S and some $n \in \mathbb{N}$, and given $j \in \mathbb{P}_n$ and $t' \in S$, we define $(\bar{t}_{-j}, t') \triangleq (t_0, \dots, t_{j-2}, t_{j-1}, t', t_{j+1}, t_{j+2}, \dots, t_{n-1}) \in S^{\mathbb{P}_n}$.
- For every $n \in \mathbb{N}$, we denote the set of permutations on \mathbb{P}_n by $\mathbb{P}_n!$.

4. PRELIMINARIES: THE CONSUMER GAME

Preparing the ground for the main results of this paper, in this section we define the congestion game among consumers, and use a novel construction³ to prove the existence of Nash equilibrium and the uniqueness of equilibrium loads (results similar in spirit to other results regarding congestion games and crowding games [Schmeidler 1973; Milchtaich 2000]), and to efficiently calculate these loads. Full proofs, as well as auxiliary results, are provided in Appendices A.1 and A.2.

Definition 4.1 (Quality-of-Service Space). For ease of presentation, we use $\mathcal{T} \triangleq [0, 1]$ as the *type space* in the consumer game (and later as the *strategy space* in the producer game). We consider lower values as indicating higher qualities of service.

For the duration of this section, fix a finite Lebesgue measure μ on \mathcal{T} , a natural $n \in \mathbb{N}$ and producer QoS levels (e.g. store distances) $\bar{t} = (t_0, \dots, t_{n-1}) \in \mathcal{T}^{\mathbb{P}_n}$. We consider the n -producers consumer game $(\mu; \bar{t}) = (\mu; t_0, \dots, t_{n-1})$, which we now define.

Definition 4.2 (Strategies). For every $d \in \mathcal{T}$, we define the set of *strategies* available to a player with type (i.e. worst acceptable QoS) d as $S_d \triangleq \{j \mid t_j \leq d\} \cup \{\neg\}$, where \neg denotes not consuming from any producer. We define $S \triangleq \cup_{d \in \mathcal{T}} S_d = \mathbb{P}_n \cup \{\neg\}$ — the set of pure strategies available to any player, and consider S as a measurable space with the σ -algebra 2^S of all of its subsets.

Definition 4.3 (Pure-Consumption Profile/Nash Equilibrium).

- (1) A *pure-consumption (strategy) profile* in the n -producers consumer game $(\mu; \bar{t})$ is a Lebesgue-measurable function $s : \mathcal{T} \rightarrow S$ s.t. $s(d) \in S_d$ for every $d \in \mathcal{T}$.
- (2) Given a pure-consumption profile s in the n -producers consumer game $(\mu; \bar{t})$, we define $\ell_j^s \triangleq \mu(s^{-1}(j))$ for every $j \in S$ — the load on producer j . (ℓ_{\neg}^s is the measure of consumers not consuming from any producer.)
- (3) A *pure-consumption Nash equilibrium* in the n -producers consumer game $(\mu; \bar{t})$ is a pure-consumption profile s s.t. for every $d \in \mathcal{T}$, both the following hold.
 - (a) $s(d) = \neg$ only if $S_d = \{\neg\}$.
 - (b) $\ell_{s(d)}^s \leq \ell_j^s$ for every $j \in S_d \setminus \{\neg\}$.⁴

We now turn to define mixed-consumption strategies. We think of such a strategy not as a probabilistic one, but rather as meaning “a certain fraction of the continuum of players with type d have one strategy, while others have other strategies”.⁵

³Note added in proof: see Gonczarowski and Tennenholtz [2014] for a generalization of this construction to arbitrary resource-selection games.

⁴As mentioned above, the results of this paper generalize also for a more general definition of the consumers’ preferences, in which each consumer consumes from producers j with minimal $f_j(\ell_j^s)$ (as opposed to minimal ℓ_j^s), where $(f_j)_{j \in \mathbb{P}_n}$ is a specification of an increasing continuous function for each producer. See Section 6 for more details.

⁵Alternatively, e.g. in the ISPs scenario given in Example 1.4, if it is possible to pay a certain ISP $\alpha \in (0, 1)$ of the price of a monthly subscription and receive an internet connection that has α of the speed of a “regular” connection, then one may also think of a mixed-consumption strategy as buying several “partial” connections, using a cumulative budget that equals the price of one regular connection (thus, if all these partial connections are purchased from ISPs with the same load, then the combined speed of all these connections would be the regular speed of a connection purchased from any one of these ISPs.)

Definition 4.4 (Mixed-Consumption Profile/Nash Equilibrium).

- (1) A *mixed-consumption (strategy) profile* in the n -producers consumer game $(\mu; \bar{t})$ is a Lebesgue-measurable function $s : \mathcal{T} \rightarrow [0, 1]^S$ s.t. $s(d) \in \Delta^{S_d}$ for every $d \in \mathcal{T}$.
- (2) Given a mixed-consumption profile s in the n -producers consumer game $(\mu; \bar{t})$, we define $\ell_j^s \triangleq \int_{\mathcal{T}} s_j d\mu$ for every $j \in S$ — the load on producer j . (ℓ_j^s is the measure of consumers not consuming from any producer in this case as well.)
- (3) A *mixed-consumption Nash equilibrium* in the n -producers consumer game $(\mu; \bar{t})$ is a mixed-consumption profile s s.t. for every $d \in \mathcal{T}$, both of the following hold.
 - (a) $\neg \in \text{supp}(s(d))$ only if $S_d = \{\neg\}$.
 - (b) $\ell_k^s \leq \ell_j^s$ for every $k \in \text{supp}(s(d))$ and $j \in S_d \setminus \{\neg\}$.⁴

THEOREM 4.5 (\exists PURE-CONSUMPTION NASH EQUILIBRIUM). *If μ is atomless, then a pure-consumption Nash equilibrium exists in the n -producers consumer game $(\mu; \bar{t})$.*

Example 4.6 (Necessity of Atomlessness Condition). Consider a nonzero measure μ concentrated entirely on the atom $d = 1 \in \mathcal{T}$. For $n > 1$, no pure-consumption Nash equilibrium exists in any induced n -producers consumer game. Indeed, in any pure-consumption profile, all consumers with type $d = 1$ would consume from the same producer, leaving another producer with a strictly-lower load of 0; as this producer is acceptable by all consumers with type $d = 1$, they would all rather deviate to it.

Definition 4.7 (Effective Type). We say that two types $d_1, d_2 \in \mathcal{T}$ are of the same *effective type* if $S_{d_1} = S_{d_2}$.

The Nash equilibrium constructed in the proof of Theorem 4.5 is asymmetric in the sense that players with the same effective type may behave differently. As we now show, this asymmetry cannot be avoided. Nonetheless, a reader who finds this asymmetry aesthetically-unpleasing may instead consider a more-symmetric, yet mixed-consumption, Nash equilibrium, which in fact exists even when μ is not atomless.

Definition 4.8 (Symmetric Strategy Profile). A strategy profile s is said to be *symmetric* if $S_{d_1} = S_{d_2} \implies s(d_1) = s(d_2)$ for every $d_1, d_2 \in \mathcal{T}$, i.e. each player's strategy depends only on the player's effective type.

Example 4.9 (Nonexistence of a Symmetric Pure-Strategy Nash Equilibrium). Consider any nonzero measure μ . For $n > 1$, if $t_j = 0$ for every $j \in \mathbb{P}_n$, then all consumers are of the same effective type. Thus, no symmetric pure-consumption equilibrium exists in the induced n -producers consumer game. Indeed, in any symmetric pure-consumption profile, since all consumers are of the same effective type, all would consume from the same producer, leaving another producer (acceptable to all) with a strictly-lower load of 0; therefore, all consumers would rather deviate to this producer.

THEOREM 4.10 (\exists SYMMETRIC MIXED-CONSUMPTION NASH EQUILIBRIUM). *A symmetric mixed-consumption Nash equilibrium exists in the n -producers consumer game $(\mu; \bar{t})$. Furthermore, there exists such an equilibrium for which the strategies can be computed in $O(n^2)$ time.*

See Fig. 1 for an illustration of the constructive proof of Theorem 4.10; as illustrated, the intuition underlying this novel construction builds upon hydraulic systems of communicating vessels (nonetheless, the proofs given in Appendix A.1 are completely formal, of course). We now show that while in general many Nash equilibria may exist in the consumer game, they result in the same load for both consumers and producers.

THEOREM 4.11 (PRODUCERS ARE INDIFFERENT BETWEEN NASH EQUILIBRIA). *$\ell_j^s = \ell_j^{s'}$ for every $j \in \mathbb{P}_n$ and every mixed-consumption Nash equilibria s, s' in $(\mu; \bar{t})$.*

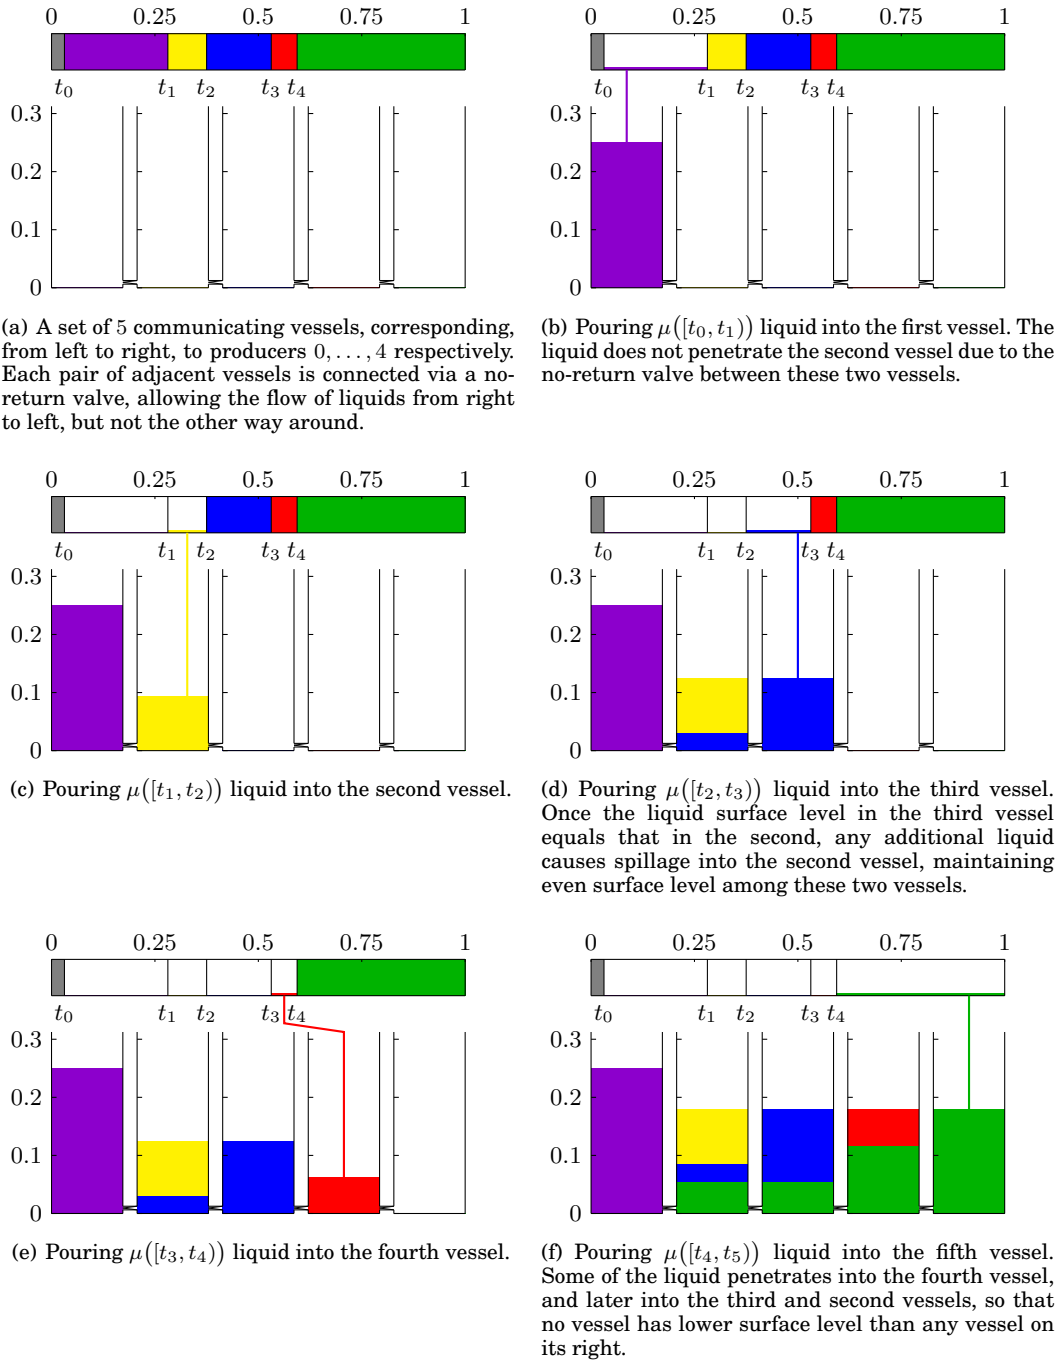


Fig. 1. Illustration of the construction in the proof of Theorem 4.10 for $n = 5$. E.g. as exactly 80% of the blue (i.e. darkest when viewed in b/w) liquid in Fig. 1(f) is in the third vessel and the remaining 20% is in the second one, the strategy for all consumer types $d \in [t_2, t_3]$ in the symmetric mixed-consumption Nash equilibrium that we construct is 0.8 consumption from producer 2 and 0.2 consumption from producer 1.

COROLLARY 4.12 (CONSUMERS ARE INDIFFERENT BETWEEN NASH EQUILIBRIA). $\ell_k^s = \ell_{k'}^s$ for every $k \in \text{supp}(s(d))$ and $k' \in \text{supp}(s'(d))$, for every $d \in \mathcal{T}$ and every mixed-consumption Nash equilibria s, s' in $(\mu; \bar{t})$.

By Theorems 4.10 and 4.11, the following is well defined.

Definition 4.13 (Producer Load). For every $j \in \mathbb{P}_n$, we define $\ell_j(\bar{t})$ to equal ℓ_j^s in any mixed-consumption Nash equilibrium s in $(\mu; \bar{t})$.

By the proof of Theorem 4.10, we obtain the following simple algorithm for directly calculating $\ell_j(\bar{t})$ for all j , without the need to first calculate consumer's strategies. While this algorithm runs in $O(n^2)$ time, i.e. has same worst-case asymptotic behaviour as explicitly computing a Nash equilibrium via Theorem 4.10 and then deducing all loads, it is considerably simpler, and also computes the loads sequentially, and so may be stopped mid-way, allowing to calculate the prefix $\ell_{\neg}(\bar{t}), \ell_0(\bar{t}), \ell_1(\bar{t}), \dots, \ell_j(\bar{t})$ of the values of $\ell(\bar{t})$ in $O(j \cdot n)$ time.

ALGORITHM 1 Direct computation of $\ell_j(\bar{t})$ for all $j \in \mathbb{P}_n$

```

1: procedure COMPUTE- $\ell(\mu; t_0, \dots, t_{n-1})$  // Assumes  $t_0 \leq t_1 \leq \dots \leq t_{n-1}$ .
2:    $t_n \leftarrow 2$  // Any value  $> 1$  will do here; assumes  $\mu$  is defined on  $[0, t_n]$  but with support  $\mathcal{T}$ .
3:    $\ell_{\neg} \leftarrow \mu([0, t_0])$ 
4:    $k \leftarrow 0$ 
5:   while  $k < n$  do
6:      $k' \leftarrow \text{Max arg Max}_{k < k' \leq n} \frac{\mu([t_k, t_{k'}])}{k' - k}$ 
7:      $\ell \leftarrow \frac{\mu([t_k, t_{k'}])}{k' - k}$ 
8:     for  $k \leq j < k'$  do
9:        $\ell_j \leftarrow \ell$ 
10:    end for
11:     $k \leftarrow k'$ 
12:  end while
13:  return  $(\ell_{\neg}, \ell_0, \dots, \ell_{n-1})$ 
14: end procedure

```

See Appendix A.2 for an analytic study of ℓ_j , formalizing some main properties thereof, which we utilize in our proofs in the following sections. In particular, we show there that for every j , $\ell_j(\bar{t})$ is nonincreasing in t_j , weakly quasiconvex in t_k for $k \neq j$, and Lipschitz (w.r.t. μ) in each coordinate with Lipschitz constant 1.

5. THE PRODUCER GAME

We now turn to the producer game, and to the main results of this paper. In this two-stage game, each producer chooses a strategy (i.e. QoS) in \mathcal{T} , and the utilities are determined according to the loads on producers in Nash equilibria in the induced consumer game. For the duration of this section, fix a natural $n \in \mathbb{N}$ and a finite Lebesgue measure μ on \mathcal{T} . Full proofs, as well as auxiliary results, are provided in Appendices A.3 and A.4.

In Section 5.1, we define a simpler version of the producer game that we define Section 5.2. While this simpler game has some trivialities that we point out, its analysis is nonetheless interesting, and the obtained results are useful when analysing the more-involved version in Section 5.2 (which is the one surveyed in the introduction).

5.1. Coarse Preferences (A Simplified Producer Game)

Definition 5.1 (Producer Game with Coarse Preferences). We define the *producer game with coarse preferences* (n, μ, \succeq_C) as the n -player game, with set of players (called *producers*) \mathbb{P}_n , in which the pure-strategy space available to each producer is \mathcal{T} , and in which for each pure-strategy profile $\bar{t} \in \mathcal{T}^{\mathbb{P}_n}$, the utility for each producer $j \in \mathbb{P}_n$ is linear in $\ell_j(\bar{t})$ (as defined in Definition 4.13).

5.1.1. Static Analysis. We begin with an analysis of domination in the producer game with coarse preferences, pointing out the trivialities in this simplified game, which will disappear in the more-involved version thereof that we analyse in Section 5.2.

Definition 5.2 (Safe Alternative; Dominant Strategy). Let t be a strategy in the game (n, μ, \succeq_C) .

- We say that t is a *safe alternative* to some strategy t' if for every strategy profile for all but one of the producers, playing t gives the remaining producer utility at least as high a utility as playing t' .
- We say that t is a *dominant strategy* if it is a safe alternative to all strategies.

THEOREM 5.3 (DOMINANT STRATEGIES). $t \in \mathcal{T}$ is a dominant strategy in (n, μ, \succeq_C) iff $\mu([0, t]) = 0$. Furthermore, each such dominant strategy guarantees a load of at least $\frac{\mu(\mathcal{T})}{n}$ on each producer playing it.

In particular, we have that every producer playing $0 \in \mathcal{T}$ constitutes a Nash equilibrium. (We emphasize that this is by far not the only Nash equilibrium — see Theorem 5.5 below.) This and other trivialities that result from domination (as well as the domination itself) disappear in Section 5.2, when we refine the order of preferences of the various producers. Before that, though, we continue to explore the consumer game with coarse preferences, obtaining results that aid our analysis of the consumer game with refined preferences in Section 5.2 below. Our next step is to not only characterize the Nash equilibrium loads (an immediate corollary of Theorem 5.3), but furthermore, show that every strategy profile inducing these loads is a Nash equilibrium.

THEOREM 5.4 (NASH EQUILIBRIUM LOADS). A pure-strategy profile $\bar{t} \in \mathcal{T}^{\mathbb{P}_n}$ constitutes a Nash equilibrium in (n, μ, \succeq_C) iff $\ell_j(\bar{t}) = \frac{\mu(\mathcal{T})}{n}$ for every $j \in \mathbb{P}_n$.

We proceed to directly characterize the strategies played in Nash equilibria, in a way that does not necessitate solving the induced consumer game.

THEOREM 5.5 (NASH EQUILIBRIUM CHARACTERIZATION). Let $t_0 \leq \dots \leq t_{n-1} \in \mathcal{T}$. The pure-strategy profile $\bar{t} \triangleq (t_1, \dots, t_{n-1})$ constitutes a Nash equilibrium in (n, μ, \succeq_C) iff $\mu([0, t_j]) \leq \frac{j}{n} \cdot \mu(\mathcal{T})$ for every $j \in \mathbb{P}_n$.

It should be emphasized that Theorem 5.5 does *not* imply that Nash equilibria are interchangeable (i.e. that the set of Nash equilibria is a Cartesian product of sets of strategies for the various producers). Consider, for example, $\mu = U(\mathcal{T})$ — the uniform measure on \mathcal{T} . In this case, by Theorem 5.5, $(0, \frac{1}{n}, \frac{2}{n}, \dots, \frac{n-1}{n})$ is a Nash equilibrium in (n, μ, \succeq_C) , and so is any permutation thereof. Nonetheless, every player playing $\frac{n-1}{n} \in \mathcal{T}$ does not constitute a Nash equilibrium. We now move on to examine the stability of the Nash equilibria in (n, μ, \succeq_C) against group deviations.

The study of stability against group deviations was initiated by Aumann [1959], who considers deviations from which all deviators gain. Recently, the CS literature considers a considerably-stronger solution concept, according to which a deviation is considered beneficial even if only some of the participants in the deviating coalition gain, as long as none of the participants lose (see e.g. [Rozenfeld and Tennenholtz 2006]).

While stability against the classical all-gaining coalitional deviation is termed *strong equilibrium*, this more-demanding concept is referred to as *super-strong equilibrium*; there are very few results showing its existence in nontrivial settings.

THEOREM 5.6 (ALL NASH EQUILIBRIA ARE SUPER STRONG). *Let $\bar{t} \in \mathcal{T}^{\mathbb{P}_n}$ be a pure-strategy Nash equilibrium in (n, μ, \succeq_C) . There exist no coalition $P \subseteq \mathbb{P}_n$ and strategies $\bar{t}' = (t'_j)_{j \in P} \in \mathcal{T}^P$ s.t. $\ell_j(\bar{t}_{-P}, \bar{t}') \geq \ell_j(\bar{t})$ for every $j \in P$, with a strict inequality for at least one producer $j \in P$.*

We conclude the static analysis of (n, μ, \succeq_C) by deducing generalizations of Theorems 5.3 to 5.6 for mixed-strategy profiles, as well as showing that no mixed-strategy Nash equilibrium exhibits any ex-post regret.

THEOREM 5.7 (MIXED STRATEGIES). *In (n, μ, \succeq_C) ,*

- (1) (DOMINANT STRATEGIES). *Let p be a mixed strategy.⁶ p is a dominant strategy iff $\mu([0, \text{Max supp}(p)]) = 0$. Furthermore, each such dominant strategy guarantees a load of at least $\frac{\mu(\mathcal{T})}{n}$ with probability 1 on each producer playing it.*
- (2) (NASH EQUILIBRIUM LOADS). *A mixed-strategy profile $\bar{p} = (p_0, \dots, p_{n-1})^7$ constitutes a Nash equilibrium iff $\ell_j(\bar{p}) = \frac{\mu(\mathcal{T})}{n}$ for every $j \in \mathbb{P}_n$ with probability 1.*
- (3) (NASH EQUILIBRIUM CHARACTERIZATION). *A mixed-strategy profile \bar{p} constitutes a Nash equilibrium iff there exists a permutation on the producers $\pi \in \mathbb{P}_n!$ s.t. $\mu([0, \text{Max supp}(p_{\pi(j)})]) \leq \frac{j}{n} \cdot \mu(\mathcal{T})$ for every $j \in \mathbb{P}_n$.*
- (4) (ALL NASH EQUILIBRIA ARE SUPER STRONG). *Let \bar{p} be a mixed-strategy Nash equilibrium. There exist no coalition $P \subseteq \mathbb{P}_n$ and mixed strategies $\bar{p}' = (p'_j)_{j \in P}$ s.t. $E[\ell_j(\bar{p}_{-P}, \bar{p}')] \geq E[\ell_j(\bar{p})]$ for every $j \in P$, with a strict inequality for at least one producer $j \in P$.*

THEOREM 5.8 (NO EX-POST REGRET IN MIXED-STRATEGY NASH EQUILIBRIA). *In any mixed-strategy Nash equilibrium in (n, μ, \succeq_C) , with probability 1 there exists no ex-post regret for any producer. In other words, a realization of a mixed-strategy Nash equilibrium is with probability 1 a pure-strategy Nash equilibrium.*

5.1.2. Dynamics. When analysing dynamics henceforth, we assume that $\mu(\mathcal{T}) > 0$. (Otherwise, by Theorem 5.3, all strategies are equivalent and so the analysis is trivial.)

Definition 5.9 (Schedule; Sequential/Simultaneous Schedule; Round).

- (1) A *schedule* is a sequence $(P_i)_{i=0}^{\infty}$ of nonempty subsets of \mathbb{P}_n , s.t. $j \in P_i$ for infinitely-many values of $i \in \mathbb{N}$, for every $j \in \mathbb{P}_n$.
- (2) A schedule $(P_i)_{i=0}^{\infty}$ is said to be *sequential* if $|P_i| = 1$ for every $i \in \mathbb{N}$.
- (3) A schedule $(P_i)_{i=0}^{\infty}$ is said to be *simultaneous* if $P_i = \mathbb{P}_n$ for every $i \in \mathbb{N}$.
- (4) Let $i_1 \leq i_2 \in \mathbb{N}$. We say that $\{i \in \mathbb{N} \mid i_1 \leq i \leq i_2\}$ constitutes a *round* (in the schedule $(P_i)_{i=0}^{\infty}$) if $\cup_{i=i_1}^{i_2} P_i = \mathbb{P}_n$. (We emphasize that this union need not be a disjoint union.)
- (5) Let $i_1 \leq i_2 \in \mathbb{N}$ and let $r \in \mathbb{N}$. We say that i_2 is *reached from i_1 in r rounds* if $r - 1$ is the largest number of pairwise-disjoint rounds into which $\{i_1, i_1 + 1, \dots, i_2 - 2\}$ can be partitioned. (Therefore, $\{i_1, i_1 + 1, \dots, i_2 - 1\}$ cannot be partitioned into r pairwise-disjoint rounds with a nonzero amount of “spare” trailing steps.)

Remark 5.10. In a simultaneous schedule, each *step* $\{i\}$ constitutes a round.

⁶We consider a *mixed strategy* to be a random variable taking values in \mathcal{T} .

⁷We emphasize that mixed-strategies of distinct producers are independent random variables.

Definition 5.11 (Weakly-/ δ -Better-/Best-Response Dynamics; Lazy Dynamics).

- (1) A *weakly-better-response dynamic* in (n, μ, \succeq_C) is a sequence $(\bar{t}_i, P_i)_{i=0}^\infty$, where $(P_i)_{i=0}^\infty$ is a schedule and $(\bar{t}^i)_{i=0}^\infty$ is a sequence of strategy profiles s.t. both of the following hold for every $i \in \mathbb{N}$.
 - For every $j \in P_i$, t_j^{i+1} is a weakly-better response than t_j^i to \bar{t}_{-j}^i (by j), i.e. $\ell_j(\bar{t}_{-j}^i, t_j^{i+1}) \geq \ell_j(\bar{t}_{-j}^i, t_j^i)$.
 - For every $j \notin P_i$, $t_j^{i+1} = t_j^i$.
- By slight abuse of notation, we sometimes write $(\bar{t}_i)_{i=0}^\infty$ to refer to $(\bar{t}_i, P_i)_{i=0}^\infty$, when the schedule is either inconsequential or clear from context.
- (2) A weakly-better-response dynamic is said to be a *best-response dynamic* if for every $i \in \mathbb{N}$ and $j \in P_i$, t_j^{i+1} is a best response to \bar{t}_{-j}^i , i.e. $t_j^{i+1} \in \arg \text{Max}_{t \in \mathcal{T}} \ell_j(\bar{t}_{-j}^i, t)$.
- (3) Let $\delta > 0$. A weakly-better-response dynamic is said to be a δ -*better-response dynamic* if for every $i \in \mathbb{N}$ and $j \in P_i$, t_j^{i+1} is either a best response to \bar{t}_{-j}^i , or a better response increasing j 's load by at least δ compared to t_j^i , i.e. $\ell_j(\bar{t}_{-j}^i, t_j^{i+1}) \geq \ell_j(\bar{t}_{-j}^i, t_j^i) + \delta$.⁸
- (4) A weakly-better-response dynamic is said to be *lazy* if for every $i \in \mathbb{N}$ and $j \in P_i$, $t_j^{i+1} = t_j^i$ whenever t_j^i is a best response to \bar{t}_{-j}^i .

Remark 5.12. In (n, μ, \succeq_C) ,

- Every best-response dynamic is a δ -better-response dynamic, for every $\delta > 0$.
- Every δ -better-response dynamic is also a δ' -better-response one, for every $0 < \delta' < \delta$.
- A weakly-better-response dynamic is a best-response dynamic iff it is a δ -better-response dynamic for $\delta = \mu(\mathcal{T})$.

Remark 5.13 (A Best Response Always Exists). Let $j \in \mathbb{P}_n$ and let $\bar{t}_{-j} \in \mathcal{T}^{\mathbb{P}_n \setminus \{j\}}$. By Theorem 5.3, a best response (by j) to \bar{t}_{-j} exists in (n, μ, \succeq_C) .

We commence with a negative result, showing that even best-response dynamics can go out of equilibrium.

Example 5.14 (Nonsequential Nonlazy Best-Response Dynamics may Go Out of Equilibrium). Let $\mu = U(\mathcal{T})$. By Theorem 5.5, the (cyclically repeating) strategy-profile sequence $(0, 0, \dots, 0)$, $(\frac{n-1}{n}, \frac{n-1}{n}, \dots, \frac{n-1}{n})$, $(0, 0, \dots, 0)$, $(\frac{n-1}{n}, \frac{n-1}{n}, \dots, \frac{n-1}{n})$, \dots constitutes a (nonlazy) simultaneous best-response dynamic in (n, μ, \succeq_C) that visits nonequilibria infinitely often.

We continue by showing that the dynamic in Example 5.14 visiting Nash equilibria infinitely often is no coincidence.

THEOREM 5.15 (δ -BETTER-RESPONSE DYNAMICS VISIT NASH EQUILIBRIA INFINITELY OFTEN). *Let $\delta > 0$ and let $(\bar{t}^i)_{i=0}^\infty$ be a δ -better-response dynamic in (n, μ, \succeq_C) . \bar{t}^i is a Nash equilibrium for infinitely-many values of i . Moreover, the first Nash equilibrium is reached (from 0) in at most $n \cdot \lceil \frac{\mu(\mathcal{T})}{\delta n} \rceil$ rounds, and from any later nonequilibrium, the next Nash equilibrium is reached in at most $(n-1) \cdot \lceil \frac{\mu(\mathcal{T})}{\delta n} \rceil$ rounds.*

Remark 5.16. In Theorem 5.15,

- if $(\bar{t}^i)_{i=0}^\infty$ is simultaneous, then “rounds” may be replaced with “steps”.
- Finer analysis of similar nature may be used to show both that $n \cdot \lceil \frac{\mu(\mathcal{T})}{\delta n} \rceil$ may be replaced with $\sum_{h=1}^n \lceil \frac{\mu(\mathcal{T})}{h\delta n} \rceil \approx \text{Max}\{\ln n \cdot \lceil \frac{\mu(\mathcal{T})}{\delta n} \rceil, n\}$, and that $(n-1) \cdot \lceil \frac{\mu(\mathcal{T})}{\delta n} \rceil$ may be

⁸Due to the continuous nature of strategies and loads, we require an improvement by at least δ , and not just any positive improvement, in order to avoid improvements *à la* Zeno’s “Race Course” paradox.

replaced with $\sum_{h=2}^n \lceil \frac{\mu(\mathcal{T})}{h\delta n} \rceil \approx \text{Max}\{(\ln n - 1) \cdot \lceil \frac{\mu(\mathcal{T})}{\delta n} \rceil, n - 1\}$. We conjecture that considerably tighter bounds (esp. for small δ) can be attained as well.

We now show that in a sense, Example 5.14 describes all the “issues” that might prevent best- and even δ -better-response dynamics from remaining in Nash equilibria.

Remark 5.17 (Lazy Better-Response Dynamics Remain in Nash Equilibrium). Once a lazy weakly-better-response dynamic reaches a Nash equilibrium, it remains constant. (Directly by definition of Nash equilibrium and laziness.)

THEOREM 5.18 (SEQUENTIAL BETTER-RESPONSE DYNAMICS REMAIN IN NASH EQUILIBRIA). *Let $(P_i)_{i=0}^\infty$ be a schedule. If $(P_i)_{i=0}^\infty$ is sequential from some point, then once a $(P_i)_{i=0}^\infty$ -scheduled weakly-better-response dynamic in (n, μ, \succeq_C) reaches a Nash equilibrium after that point, it never visits a nonequilibrium afterward.*

COROLLARY 5.19 (SEQUENTIAL/LAZY δ -BETTER-RESPONSE DYNAMICS REACH NASH EQUILIBRIA AND REMAIN). *For every $\delta > 0$, every sequential or lazy δ -better-response dynamic in (n, μ, \succeq_C) reaches a Nash equilibrium in a finite number of steps, and never visits a nonequilibrium after that point.*

PROPOSITION 5.20 (EVERY NONSEQUENTIAL SCHEDULE AND INITIAL PROFILE HAVE NONLAZY BEST-RESPONSE DYNAMICS THAT GO OUT OF EQUILIBRIUM). *If μ has no atom measuring $\frac{n-1}{n} \cdot \mu(\mathcal{T})$ or more and no tail of $(P_i)_{i=0}^\infty$ is sequential, then for every pure-strategy profile t^0 there exists a nonlazy best-response dynamic in (n, μ, \succeq_C) that is scheduled by $(P_i)_{i=0}^\infty$, starts at t^0 and visits nonequilibria infinitely often.*

Remark 5.21. Analogues of Remark 5.17, Theorem 5.18, and Proposition 5.20 hold for mixed-strategy dynamics as well.

As every best-response dynamic is a δ -better-response one for $\delta = \mu(\mathcal{T})$, we conclude from Theorem 5.15 that such a dynamic reaches a Nash equilibrium in at most $n \cdot \lceil \frac{\mu(\mathcal{T})}{\mu(\mathcal{T})n} \rceil = n$ rounds and afterward always “re-reaches” a Nash equilibrium in at most $(n - 1) \cdot \lceil \frac{\mu(\mathcal{T})}{\mu(\mathcal{T})n} \rceil = n - 1$ rounds. By applying some finer analysis, we can slightly improve this bound, and show that the new bound is tight.

THEOREM 5.22 (BEST-RESPONSE TIME-TO-EQUILIBRIUM AND TIME BETWEEN EQUILIBRIA). *Every best-response dynamic in (n, μ, \succeq_C) reaches a Nash equilibrium in at most $n - 1$ rounds. Furthermore, if $n > 2$, then from any later nonequilibrium, the next Nash equilibrium is reached in at most $n - 2$ rounds.*

Remark 5.23. In Theorem 5.22, as in Theorem 5.15, if the dynamic in question is simultaneous, then “rounds” may be replaced with “steps”.

Example 5.24 (Tightness of Theorem 5.22). Let $\mu = U(\mathcal{T})$. The following is a (nonlazy) simultaneous best-response dynamic in (n, μ, \succeq_C) , in which i) no two consecutive strategy profiles are both Nash equilibria, ii) the first Nash equilibrium is reached in precisely $n - 1$ rounds (steps), and iii) from any nonequilibrium that follows a Nash equilibrium, the next Nash equilibrium is reached in precisely $n - 2$ rounds (steps): $(1, 0, 0, \dots, 0)$, $(\frac{n-1}{n}, \frac{n-2}{n-1}, 0, 0, \dots, 0)$, $(\frac{n-2}{n}, \frac{n-2}{n}, \frac{n-3}{n-1}, 0, 0, \dots, 0)$, \dots , $(\frac{3}{n}, \frac{3}{n}, \dots, \frac{3}{n}, \frac{2}{n-1}, 0, 0)$, $(\frac{2}{n}, \frac{2}{n}, \dots, \frac{2}{n}, \frac{1}{n-1}, 0)$, $(\frac{1}{n}, \frac{1}{n}, \dots, \frac{1}{n}, 0)$ (first Nash equilibrium), $(\frac{n-1}{n}, \frac{n-2}{n-1}, 0, 0, \dots, 0)$, $(\frac{n-2}{n}, \frac{n-2}{n}, \frac{n-3}{n-1}, 0, 0, \dots, 0)$, \dots (cyclically repeating).

To summarize Theorems 5.18 and 5.22, Example 5.24, and Remark 5.17:

COROLLARY 5.25 (SEQUENTIAL/LAZY BEST-RESPONSE DYNAMICS REACH NASH EQUILIBRIA FAST AND REMAIN). *Every sequential or lazy best-response dynamic in*

(n, μ, \succeq_C) reaches a Nash equilibrium in at most a tight bound of $n - 1$ rounds, and never visits a nonequilibrium after that point.

5.2. Fine Preferences

Definition 5.26 (Producer Game with Fine Preferences). We define the *producer game with fine preferences* (n, μ, \succeq_F) as the n -player game, with set of players (called *producers*) \mathbb{P}_n , in which the pure-strategy space available to each producer is \mathcal{T} , and in which for each pure-strategy profile $t \in \mathcal{T}^{\mathbb{P}_n}$, the utility for each producer $j \in \mathbb{P}_n$ is linear in $\ell_j(t)$ (as defined in Definition 4.13), with tie breaking (i.e. infinitesimal improvement) in favour of larger values of t_j over smaller ones.

5.2.1. Static Analysis. We define safe alternatives and dominant strategies in (n, μ, \succeq_F) as in Definition 5.2, only w.r.t. fine preferences. The following proposition shows that the tie-breaking refinement of the producers' preferences into "fine preferences" indeed successfully removes the triviality captured by Theorem 5.3, in a strong sense.

PROPOSITION 5.27 ((NO) DOMINANT AND (FEW) DOMINATED STRATEGIES). *If μ has no atom measuring $\frac{n-1}{n} \cdot \mu(\mathcal{T})$ or more,⁹ then no strategies are dominant in (n, μ, \succeq_F) . Moreover, at least $\frac{n-1}{n}$ of the strategies in \mathcal{T} (as measured by μ) have no safe alternatives (other than themselves).*

We now formally conclude the results captured informally in Theorem 1.2:

THEOREM 5.28 ($\exists!$ NASH EQUILIBRIUM, AND IT IS SUPER STRONG¹⁰). *A unique (up to permutations) pure-strategy Nash equilibrium exists in (n, μ, \succeq_F) . The sorted Nash-equilibrium strategies $t_0 \leq \dots \leq t_{n-1} \in \mathcal{T}$ are $t_j \triangleq \text{Max}\{t \in \mathcal{T} \mid \mu([0, t]) \leq \frac{j}{n} \cdot \mu(\mathcal{T})\}$ for every $j \in \mathbb{P}_n$. The load on each producer in this equilibrium is $\frac{\mu(\mathcal{T})}{n}$. Furthermore, this equilibrium is super strong.*

COROLLARY 5.29 (NASH EQUILIBRIUM CHARACTERIZATION — SPECIAL CASE). *If the CDF of μ is continuous (i.e. μ is atomless) and strictly increasing, then for every $j \in \mathbb{P}_n$, the j^{th} sorted Nash-equilibrium strategy, t_j , is the unique strategy satisfying $\mu([0, t_j]) = \frac{j}{n} \cdot \mu(\mathcal{T})$.*

PROPOSITION 5.30 (NASH EQUILIBRIA ARE IN PURE STRATEGIES). *Every mixed-strategy Nash equilibria in (n, μ, \succeq_F) is in fact in pure strategies (and is thus given by Theorem 5.28 / Corollary 5.29).*

If μ is atomless, then in the Nash equilibrium defined in Theorem 5.28 and Corollary 5.29, almost all (i.e. except for maybe an amount of measure zero) of the $1/n$ of consumers (as measured by μ) with numerically-smallest types consume from producer 0, whose chosen strategy is the numerically-largest one that accommodates almost all of this $1/n$; almost all of the next $1/n$ of consumers consume from producer 1, whose chosen strategy is the numerically-largest one that accommodates almost all of this $1/n$, and so forth. Essentially, the market is allocated (i.e. split) among the various producers based on consumer types, and each producer chooses the numerically-largest strategy that accommodates almost all of its allocated market share. We conclude the static analysis of (n, μ, \succeq_F) by formalizing these results.

⁹The triviality in case of a very large atom should be compared to the triviality captured under the exact condition in Proposition 5.20. Indeed, both trivialities are possible exactly iff there exists $t \in \mathcal{T}$ s.t. $\mu([0, t]) = 0$ and $\mu([t', 1]) < \frac{\mu(\mathcal{T})}{n}$ for every $t' > t$.

¹⁰See Theorem 5.6 in Section 5.1 above, as well as the preceding discussion, for the definition of super-strong equilibrium, as well as a discussion regarding various group-deviation concepts.

THEOREM 5.31 (MARKET ALLOCATION). *Let $t_0 \leq \dots \leq t_{n-1} \in \mathcal{T}$ s.t. \bar{t} constitutes a Nash equilibrium in (n, μ, \succeq_F) , and let s be a mixed-consumption Nash equilibrium in the induced consumer game $(\mu; \bar{t})$. If μ is atomless, then for every $j \in \mathbb{P}_n$, $s_j(d) = 1$ for almost all (w.r.t. μ) consumer types $d \in \mathcal{T}$ s.t. $\mu([0, d]) \in (\frac{j}{n} \cdot \mu(\mathcal{T}), \frac{j+1}{n} \cdot \mu(\mathcal{T}))$.*

Remark 5.32 (Producer Strategies are Chosen according to the Market Allocation). By Theorem 5.28, if μ is atomless, then the j^{th} sorted Nash-equilibrium strategy, t_j , is the numerically-largest strategy (i.e. lowest QoS) acceptable by almost all consumer types $d \in \mathcal{T}$ s.t. $\mu([0, d]) \in (\frac{j}{n} \cdot \mu(\mathcal{T}), \frac{j+1}{n} \cdot \mu(\mathcal{T}))$.

5.2.2. Dynamics. We define weakly- δ -better/best-response dynamics in (n, μ, \succeq_F) as in Definition 5.11, only with best responses defined w.r.t. fine preferences. In particular, the definition of improvement by at least δ remains unchanged (i.e. it is defined solely w.r.t. the load). The last part of Remark 5.12 thus becomes:

Remark 5.33. In (n, μ, \succeq_F) , a weakly-better-response dynamic is a best-response dynamic iff it is a δ -better-response dynamic for some (equivalently, for all) $\delta > \mu(\mathcal{T})$.

We start by noting that best responses always exist — an observation that for general μ is considerably less trivial w.r.t. fine preferences than w.r.t. coarse ones.

PROPOSITION 5.34 (A UNIQUE BEST RESPONSE ALWAYS EXISTS). *Let $j \in \mathbb{P}_n$ and let $\bar{t}_{-j} \in \mathcal{T}^{\mathbb{P}_n \setminus \{j\}}$. A unique best response (by j) to \bar{t}_{-j} exists in (n, μ, \succeq_F) .*

We give two proofs for Proposition 5.34: the first — quite-concise, and the second, while requiring more involved arguments, is constructive in the sense that in contrast to the first, it presents the best response in the form $\text{Max}\{t \in \mathcal{T} \mid \mu([0, t]) \leq m\}$, for m that can be *explicitly* calculated.

As in (n, μ, \succeq_F) no producer is ever indifferent between two strategies, all weakly-better-response dynamics in this game are lazy; therefore, if such a dynamic reaches a Nash equilibrium, it remains constant from that point on. We also note that Example 5.14 is also an example of a lazy best-response dynamic in (n, μ, \succeq_F) that never reaches a Nash equilibrium. Moreover, as best responses in (n, μ, \succeq_F) are unique, and as the strategies in each Nash equilibrium are distinct if μ is atomless, we have that no simultaneous best-response dynamic starting from a strategy profile with two or more identical strategies ever reaches a Nash equilibrium. It should be noted that this is not a boundary phenomenon; for example, if $\mu([0, t_j^0]) > 0$ for all $j \in \mathbb{P}_n$, then the best responses of all producers are identical (see Corollary A.29 in Appendix A.4). Many more such examples may be constructed. We now show that these phenomena are all avoided by sequential dynamics.

COROLLARY 5.35. *Theorems 5.15 and 5.22 hold also regarding reaching a Nash equilibrium w.r.t. (n, μ, \succeq_C) by dynamics in the game (n, μ, \succeq_F) .*

THEOREM 5.36 (SEQUENTIAL δ -BETTER-RESPONSE DYNAMICS CONVERGE FROM COARSE-PREFERENCES NASH EQUILIBRIUM). *If $(P_i)_{i=0}^\infty$ is sequential from some point, then for every $\delta > 0$, at most one round after a δ -better-response dynamic in (n, μ, \succeq_F) reaches a Nash equilibrium w.r.t. (n, μ, \succeq_C) after that point, it reaches a Nash equilibrium w.r.t. (n, μ, \succeq_F) , and remains constant from that point onward.*

We hence formally conclude the results captured informally in Theorem 1.3:

COROLLARY 5.37 (SEQUENTIAL δ -BETTER-RESPONSE DYNAMICS CONVERGE). *For every $\delta > 0$, every sequential δ -better-response dynamic in (n, μ, \succeq_F) reaches a Nash equilibrium in a finite number of steps, and remains constant from that point onward.*

COROLLARY 5.38 (SEQUENTIAL BEST-RESPONSE DYNAMICS CONVERGE FAST). *Every sequential or lazy best-response dynamic in (n, μ, \succeq_F) reaches a Nash equilibrium in at most n rounds, and never visits a nonequilibrium after that point.*

We conjecture than an even-tighter bound on convergence time than in Corollaries 5.35 and 5.38 is attainable.

Corollaries 5.37 and 5.38 show that the seemingly-cooperative market allocation (split) among the various producers shown above to be exhibited in every Nash equilibrium in (n, μ, \succeq_F) arises as the unique possible outcome, not as a result of anticompetitive practices, but rather as a result of noncooperative dynamics, each producer only looking to myopically maximize its preferences at every step; as the best response to any strategy profile is unique, no signalling or any other collusive or cooperative “trick” whatsoever is used in order to reach and maintain this market allocation.

6. HETEROGENEOUS PRODUCTS

We have so far assumed that each consumer wishes to consume from producers with least loads. Recall Example 1.4 from the introduction, regarding ISPs and customers; in this example, we may imagine that some ISPs may have different total bandwidth than others, while some other ISPs may purchase more total bandwidth as their subscriber pool grows. In such a scenario, in order to surf with greatest speed, each consumer would no longer like to consume from a producer with least ℓ_j (i.e. with as few subscribers as possible), but would rather consume from a producer with least $f_j(\ell_j)$, where f_j is an increasing continuous function for every $j \in \mathbb{P}_n$, possibly differing between producers. The results in this paper lend to generalization also to such a scenario via similar methods, with only quantitative rather than qualitative changes; notably, the unique market-share allocation in both fine- and coarse-preferences Nash equilibria among producers is generally no longer the allocation of $1/n$ of the market to each of the producers. E.g. Theorems 5.3 to 5.5 thus become:

THEOREM 6.1 (HETEROGENEOUS PRODUCTS — COARSE PREFERENCES). *There exist amounts $\tilde{\ell}_0, \tilde{\ell}_1, \dots, \tilde{\ell}_{n-1} \in [0, \mu(\mathcal{T})]$ (for homogeneous products, $\tilde{\ell}_j = \frac{\mu(\mathcal{T})}{n}$ for every $j \in \mathbb{P}_n$), s.t. all of the following hold.*

- (1) (DOMINANT STRATEGIES). *Each dominant strategy in (n, μ, \succeq_C) (the characterization of such strategies is unchanged from that given in the first part of Theorem 5.3), when played by a producer $j \in \mathbb{P}_n$, guarantees a load of at least $\tilde{\ell}_j$ on this producer.*
- (2) (NASH EQUILIBRIUM LOADS). *A pure-strategy profile $\bar{t} \in \mathcal{T}^{\mathbb{P}_n}$ constitutes a Nash equilibrium in (n, μ, \succeq_C) iff $\ell_j(\bar{t}) = \tilde{\ell}_j$ for every $j \in \mathbb{P}_n$.*
- (3) (NASH EQUILIBRIUM CHARACTERIZATION). *Let \bar{t} be a pure-strategy profile and let $\pi \in \mathbb{P}_n!$ be a permutation s.t. $t_{\pi(0)} \leq t_{\pi(1)} \leq \dots \leq t_{\pi(n-1)}$. \bar{t} constitutes a Nash equilibrium in (n, μ, \succeq_C) iff $\mu([0, t_{\pi(j)}]) \leq \sum_{k=0}^{j-1} \tilde{\ell}_{\pi(k)}$ for every $j \in \mathbb{P}_n$.*

Consequently, Theorem 5.28 and Corollary 5.29 become:

THEOREM 6.2 (HETEROGENEOUS PRODUCTS — FINE PREFERENCES).

- (1) ($\exists!$ NASH EQUILIBRIUM, AND IT IS SUPER STRONG). *Let $\pi \in \mathbb{P}_n!$ be a permutation s.t. there do not exist $j < k \in \mathbb{P}_n$ s.t. $\tilde{\ell}_{\pi(j)} = 0$ while $\tilde{\ell}_{\pi(k)} \neq 0$. A unique pure-strategy Nash equilibrium s.t. $t_{\pi(0)} \leq \dots \leq t_{\pi(n-1)}$ exists in (n, μ, \succeq_F) . The strategies of this equilibrium are given by $t_{\pi(j)} \triangleq \text{Max}\{t \in \mathcal{T} \mid \mu([0, t]) \leq \sum_{k=0}^{j-1} \tilde{\ell}_{\pi(k)}\}$ for every $j \in \mathbb{P}_n$. The load on each producer $j \in \mathbb{P}_n$ in this Nash equilibrium is $\tilde{\ell}_j$. Furthermore, this equilibrium is super strong. No other Nash equilibria exist in (n, μ, \succeq_F) .*




Fig. 2. A system of 5 one-way communicating vessels, corresponding to 5 heterogeneous ISPs (see Example 1.4) with the following characteristics, from left to right (i.e. from lowest latency/highest QoS to highest latency/lowest QoS): A “normal” ISP, an ISP with half the total bandwidth of a “normal” one, an ISP whose total bandwidth somewhat increases with its number of subscribers, an ISP whose total bandwidth somewhat decreases with its number of subscribers, and a “normal” ISP who buys additional bandwidth if needed, so that the bandwidth for a single subscriber never drops below some threshold. (After the surface of the liquid in the fifth vessel reaches the tube connecting this vessel to the container on its right, which we also consider as part of the fifth vessel, any additional liquid poured into this vessel accumulates in the container on the right; assume that this container is large enough so as to never fill up.) We emphasize that the technical modifications to Lemma A.2 to accommodate any collection of increasing continuous functions $(f_j)_{j \in \mathbb{P}_n}$ are straightforward and do not require defining any shapes for any vessels — this is done purely to convey intuition. (We require that the functions be strictly increasing for simplicity, however these results still hold if one of them is merely nondecreasing, e.g. as in the scenario depicted in the figure; however, if more than one of these functions is not strictly increasing, e.g. if a sixth vessel identical to the fifth one is added in this figure, then Theorem 4.11 may no longer hold.) The producer-equilibrium loads $\tilde{\ell}_0, \dots, \tilde{\ell}_{n-1}$ can be found by pouring the entire $\mu(\mathcal{T})$ of liquid into the rightmost vessel (i.e. computing the loads when each producer j ’s strategy is the $\tilde{\ell}_j$ -guaranteeing strategy $0 \in \mathcal{T}$), or, equivalently, by simply removing the one-way valves (i.e. permitting liquid flow in both directions) and pouring $\mu(\mathcal{T})$ liquid into the system (observe that either way, if all vessels are of the same shape, then we indeed obtain $\tilde{\ell}_j = \frac{\mu(\mathcal{T})}{n}$ for every $j \in \mathbb{P}_n$, as in the previous sections); a similar “two-way” calculation among subsets of vessels generalizes Algorithm 1 to this scenario.

- (2) (NASH EQUILIBRIUM CHARACTERIZATION — SPECIAL CASE). *If the CDF of μ is continuous (i.e. μ is atomless) and strictly increasing, then for every $j \in \mathbb{P}_n$, in the Nash equilibrium corresponding to a permutation $\pi \in \mathbb{P}_n!$ with the above properties, $t_{\pi(j)}$ is the unique strategy satisfying $\mu([0, t_{\pi(j)}]) = \sum_{k=0}^{j-1} \tilde{\ell}_{\pi(k)}$.*

The remainder of the results of this paper, including those regarding dynamics, readily generalize to this scenario as well. So, we once again have that in a Nash equilibrium in (n, μ, \succeq_F) , the market is allocated (split) among producers based on consumer types; if μ is atomless and $t_{\pi(0)} \leq \dots \leq t_{\pi(n-1)}$, then almost all of the $\tilde{\ell}_{\pi(0)}$ consumers with numerically-smallest types consume from producer $\pi(0)$ (who chooses the largest strategy acceptable by almost all of them), almost all of the next $\tilde{\ell}_{\pi(1)}$ consumers consume from producer $\pi(1)$ (who chooses the largest strategy acceptable by almost all of them), and so forth. See Fig. 2 for an illustration regarding the adaptation of the results from Section 4 to this generalized model, and the calculation of $\tilde{\ell}_0, \dots, \tilde{\ell}_{n-1}$.¹¹

7. THE FORMATION OF MAIN STREET

We conclude the body of this paper with an aesthetically-appealing corollary. Let us consider a two-goods facility-location problem such as that in the introduction. (The results of this section readily generalize also to the case of more than two good types.) Formally, we have $n_1 \in \mathbb{N}$ producers of the first good (e.g. wine) and $n_2 \in \mathbb{N}$ producers of the second good (e.g. olive oil). The strategy of each producer is a point on the plane; a pure-consumption strategy of a consumer with type $d \in \mathcal{T}$ is a pair $(j, k) \in \mathbb{P}_{n_1} \times \mathbb{P}_{n_2}$,

¹¹Note added in proof: see Gonczarowski and Tennenholtz [2014] for a formalization, as well as (as mentioned above) a generalization to arbitrary resource-selection games.

denoting consumption of the first good from producer j of this good, and of the second good — from producer k of that good; each consumer would like to minimize $f_j^1(\ell_j) + f_k^2(\ell_k)$ ¹² (e.g. the sum of the quotients of the quality and circulation for each good), subject to the constraint that the circumference of the triangle whose vertices are the origin and the locations (strategies) of producer j of good 1 and of producer k of good 2 does not exceed $2d$ (the density of consumer types, as given by μ , remains unchanged). Each producer would like to first and foremost maximize its number of consumers, and only as a tie-breaker, maximize the norm of its strategy (i.e. its distance from the origin). Under these conditions, roughly speaking, each producer would like to be located so that visiting it would never be too much of a detour on the way from the origin to a producer of the other good. Formally, we obtain: (See Appendix A.5 for a proof, as well as a discussion regarding the necessity of the conditions below.)

THEOREM 7.1 (THE UNIQUE SUPER-STRONG EQUILIBRIUM IS A MAIN STREET ORIGINATING FROM THE ORIGIN). *Let $\tilde{\ell}_0^1, \dots, \tilde{\ell}_{n_1-1}^1$ be the producer-equilibrium loads when only the first good is on the market (i.e. as defined in Section 6 when the only producers are the n_1 producers of good 1) and let $\tilde{\ell}_0^2, \dots, \tilde{\ell}_{n_2-1}^2$ be the producer-equilibrium loads when only the second good is on the market. If no nonempty proper subset of the former loads and no nonempty proper subset of the latter loads have the same sum, and if $\tilde{\ell}_j^g > 0$ for all g, j , then a producer strategy profile is a super-strong equilibrium iff the strategies of all producers of both products are on the same ray from the origin, with distances from the origin as in Theorem 6.2 (when computed separately for each good).*

While most readers are likely to consider the formation of a main street as a fairly-natural phenomenon due to its abundance in many cities, some readers may find it somewhat less natural for this main street, as deduced in Theorem 7.1, to originate from the city centre (named Metropolis Central in the introduction), rather than having the city centre in its middle. Such readers may compare this with the structure of many old European towns, at the heart of which lies the old stone-cobbled main street, on one end of which (as opposed to at the middle of which) lies the main town church.

ACKNOWLEDGEMENTS

This work was supported in part by ISF grant 230/10 and by the European Research Council under the European Community's Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement no. [249159]. We would like to thank Sergiu Hart for pointing out an intuitive resemblance between our communicating-vessels analogy and that of Kaminsky [2000].

REFERENCES

- ASHLAGI, I., LUCIER, B., AND TENNENHOLTZ, M. 2013. Equilibria of online scheduling algorithms. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence*. 67–73.
- ASHLAGI, I., TENNENHOLTZ, M., AND ZOHAR, A. 2010. Competing schedulers. In *Proceedings of the 24th AAAI Conference on Artificial Intelligence*. 691–696.
- AUMANN, R. J. 1959. Acceptable points in general cooperative n -person games. In *Contributions to the Theory of Games, Volume IV*, A. W. Tucker and R. D. Luce, Eds. Annals of Mathematics Studies Series, vol. 40. Princeton University Press, Princeton, UK, 287–324.
- GONCZAROWSKI, Y. A. AND TENNENHOLTZ, M. 2014. Physical computation of Nash equilibrium: ‘hydraulic’ resource selection. In preparation.

¹²This sum may be replaced with any increasing continuous function of f_j^1, f_k^2 , e.g. their weighted average.

- HOLZMAN, R. 2003. Network structure and strong equilibrium in route selection games. *Mathematical Social Sciences* 46, 2, 193–205.
- HOLZMAN, R. AND LAW-YONE, N. 1997. Strong equilibrium in congestion games. *Games and Economic Behavior* 21, 1, 85–101.
- HOTELLING, H. 1929. Stability in competition. *Economic Journal* 39, 153, 41–57.
- KAMINSKY, M. M. 2000. ‘Hydraulic’ rationing. *Mathematical Social Sciences* 40, 2, 131–155.
- MCAFEE, R. 1993. Mechanism design by competing sellers. *Econometrica: Journal of the Econometric Society* 61, 6, 1281–1312.
- MILCHTAICH, I. 2000. Generic uniqueness of equilibrium in large crowding games. *Mathematics of Operations Research* 25, 3, 349–364.
- MONDERER, D. AND SHAPLEY, L. S. 1996. Potential games. *Games and Economic Behavior* 14, 1, 124–143.
- ROSENTHAL, R. W. 1973. A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory* 2, 1, 65–67.
- ROUGHGARDEN, T. AND TARDOS, É. 2002. How bad is selfish routing? *Journal of the ACM* 49, 2, 236–259.
- ROZENFELD, O. AND TENNENHOLTZ, M. 2006. Strong and correlated strong equilibria in monotone congestion games. In *Proceedings of the 2nd international Workshop on Internet & Network Economics (WINE)*. 74–86.
- SCHMEIDLER, D. 1973. Equilibrium points of nonatomic games. *Journal of Statistical Physics* 7, 4, 295–300.

APPENDIX

A. PROOFS AND ADDITIONAL RESULTS

A.1. Proofs and Additional Results for Section 4

We commence with a few lemmas used in the proofs of Theorems 4.5 and 4.10.

LEMMA A.1 (LOAD IS NONINCREASING IN STRATEGY). *Under the definitions of Section 4, if $t_0 \leq t_1 \leq \dots \leq t_{n-1}$, then for every mixed-consumption Nash equilibrium s in the n -producers consumer game $(\mu; \bar{t})$, we have $\ell_0^s \geq \ell_1^s \geq \dots \geq \ell_{n-1}^s$.*

PROOF. Let $j \in \{0, \dots, n-2\}$. If $\ell_{j+1}^s = 0$, then $\ell_j^s \geq \ell_{j+1}^s$. Assume, therefore, that $\ell_{j+1}^s > 0$. Hence, there exists $d \geq t_{j+1}$ s.t. $s_{j+1}(d) > 0$. By definition of s , and as $d \geq t_{j+1} \geq t_j$, we thus have $\ell_{j+1}^s \leq \ell_j^s$, as required. \square

As mentioned above, the construction in the proof of Theorems 4.5 and 4.10 is illustrated in Fig. 1. In the context of that figure, the following lemma can be thought of as answering the following question: if the amount of liquid in each vessel $j \in \mathbb{P}_n$ is ℓ_j , by how much would the liquid in each vessel rise if we pour an additional amount m of liquid into vessel $n-1$? (The rise in the amount of liquid in vessel j is given by p_j .)

LEMMA A.2. *Let $\ell_0 \geq \ell_1 \geq \dots \geq \ell_{n-1}$ be a finite nonincreasing sequence in \mathbb{R}_{\geq} . For every $m \in \mathbb{R}_{\geq}$, there exists $p \in [0, m]^{\mathbb{P}_n}$, which may be computed in $O(n)$ time, s.t. all of the following hold.*

- (1) $\sum_{j=0}^{n-1} p_j = m$.
- (2) $\ell_0 + p_0 \geq \ell_1 + p_1 \geq \dots \geq \ell_{n-1} + p_{n-1}$.
- (3) $\ell_k + p_k = \min_{j \in \mathbb{P}_n} \{\ell_j + p_j\}$ for every $k \in \mathbb{P}_n$ s.t. $p_k > 0$.

PROOF. We iteratively define a sequence $p^n \leq p^{n-1} \leq \dots \leq p^0 \in [0, m]^{\mathbb{P}_n}$ s.t. the following hold for every $i \in \{0, \dots, n\}$.

- (1) $p_j^i = 0$ for every $j < i$.
- (2) $\sum_{j=0}^{n-1} p_j^i \leq m$, with equality when $i = 0$.
- (3) There exists $h_i \in \mathbb{R}_{\geq}$ s.t. all of the following hold.
 - If $\sum_{j=0}^{n-1} p_j^{i+1} < m$, then $\ell_j + p_j^i = h_i$ for every $j \geq i$,
 - If $\sum_{j=0}^{n-1} p_j^i < m$, then $\ell_{i-1} + p_{i-1}^i = h_i$ as well.
 - $\ell_j \geq h_i$ for every $j < i$.

In the setting of Fig. 1, p^n describes the rise of liquid before we begin pouring the additional amount m , while for every $i \in \mathbb{P}_n$, p^i describes the rise of liquid at the last instant during the pouring process, in which no water has risen except in vessels $i, i+1, \dots, n-1$. (This can be either the final rise in liquid if $i = 0$ or if the final rise does not involve a change in the amount of liquid in vessels $j < i$, or alternatively the rise in liquid just before the liquid in vessel $i-1$ begins to rise.)

For the base case, we define $p^n \equiv 0$, and all parts trivially hold (with $h_n \triangleq \ell_{n-1}$). For the construction step, let $i \in \{0, \dots, n-1\}$ and assume that p^{i+1} has been defined. Let $c \triangleq \sum_{j=0}^{n-1} p_j^{i+1}$. By Property 2 for $i+1$, $c \leq m$. If $i = 0$, then we define $r \triangleq m - c \geq 0$; otherwise, we define $r \triangleq \min\{(n-i) \cdot (\ell_{i-1} - h_{i+1}), m - c\}$, and by Property 3, $r \geq 0$ in this case as well. We define $p_j^i \triangleq 0$ for every $j < i$ (and so Property 1 holds for i), and $p_j^i \triangleq p_j^{i+1} + \frac{r}{n-i} \geq p_j^{i+1}$ for every $j \geq i$. Property 2 holds for i as $\sum_{j=0}^{n-1} p_j^i = \sum_{j=0}^{n-1} p_j^{i+1} + r = c + r \leq m$, with equality when $i = 0$. Finally, we show that Property 3 holds for i , with $h_i \triangleq h_{i+1} + \frac{r}{n-i}$. If $\sum_{j=0}^{n-1} p_j^{i+1} < m$, then as $p^{i+2} \leq p^{i+1}$, we have that

($i + 1 = n$ or) $\sum_{j=0}^{n-1} p_j^{i+2} < m$ as well. Therefore, by Property 3 for $i + 1$, we have for every $j \geq i$ that $\ell_j + p_j^i = \ell_j + p_j^{i+1} + \frac{r}{n-i} = h_{i+1} + \frac{r}{n-i} = h_i$. If $\sum_{j=0}^{n-1} p_j^i < m$, then $r < m - c$ and so by definition, $r = (n - i) \cdot (\ell_{i-1} - h_{i+1})$. Therefore, $h_i = h_{i+1} + \frac{r}{n-i} = \ell_{i-1} = \ell_{i-1} + p_{i-1}^i$. Finally, for every $j < i$, by ℓ nonincreasing and by definition of r we have $\ell_j \geq \ell_{i-1} \geq h_{i+1} + \frac{r}{n-i} = h_i$, and the proof of the construction is complete.

Let $\tilde{i} \in \{0, \dots, n\}$ be largest s.t. $\sum_{j=0}^{n-1} p_j^{\tilde{i}} = m$; \tilde{i} is well defined by Property 2 for $i = 0$. We now show that $p \triangleq p^{\tilde{i}}$ meets the conditions of the lemma. By definition, $\sum_{j=0}^{n-1} p_j = m$.

Let $j \in \mathbb{P}_n \setminus \{n-1\}$. If $j < \tilde{i} - 1$, then by Property 1 for $i = \tilde{i}$, we have $\ell_j + p_j = \ell_j \geq \ell_{j+1} = \ell_{j+1} + p_{j+1}$. If $j = \tilde{i} - 1$, then by Properties 1 and 3 for $i = \tilde{i}$ and by definition of \tilde{i} , we have $\ell_j + p_j = \ell_j = \ell_{\tilde{i}-1} \geq h_{\tilde{i}} = \ell_{\tilde{i}} + p_{\tilde{i}} = \ell_{j+1} + p_{j+1}$. Otherwise, i.e. if $j > \tilde{i} - 1$, then by Property 3 for $i = \tilde{i}$ and by definition of \tilde{i} , we have $\ell_j + p_j = h_{\tilde{i}} = \ell_{j+1} + p_{j+1}$.

We conclude that $\min_{j \in \mathbb{P}_n} \{\ell_j + p_j\} = \ell_{n-1} + p_{n-1}$. For every $k \in \mathbb{P}_n$ s.t. $p_k > 0$, by Property 1 for $i = \tilde{i}$ we have $k \geq \tilde{i}$. Therefore, by Property 3 for $i = \tilde{i}$ and by definition of \tilde{i} , we have $\ell_k + p_k = h_{\tilde{i}} = \ell_{n-1} + p_{n-1} = \min_{j \in \mathbb{P}_n} \{\ell_j + p_j\}$, as required.

Finally, although it may seem in first glance that $O(n^2)$ time may be required to compute p , we note that the sequence $(h_i)_{i=0}^n$ can be computed in $O(n)$ time, that from this sequence \tilde{i} can be deduced in $O(n)$ time as the largest s.t. $h_{\tilde{i}} = h_0$, and that from both, p can be calculated in $O(n)$ time: $p_j = 0$ for $j < \tilde{i}$ by Property 1 for $i = \tilde{i}$, while $p_j = h_{\tilde{i}} - \ell_j$ for $j \geq \tilde{i}$ by Property 3 for $i = \tilde{i}$. \square

We now prove Theorem 4.10, and then deduce Theorem 4.5 therefrom. Alternatively, Theorem 4.5 can also be proven directly from Lemma A.2, similarly to the proof of Theorem 4.10.

Definition A.3. For a finite Lebesgue measure μ on \mathcal{T} and a Lebesgue-measurable set $E \subseteq \mathcal{T}$, we denote by $\mu|_{\cap E}$ the finite Lebesgue measure on \mathcal{T} given by $\mu|_{\cap E}(A) \triangleq \mu(A \cap E)$.

PROOF OF THEOREM 4.10. We prove the claim by induction on n . (Recall that the construction underlying this proof is illustrated by Fig. 1; also recall the explanation preceding the statement of Lemma A.2 regarding the meaning of that lemma in the context of that figure.)

Base ($n = 0$): In this case, $S = \{\neg\}$, and so $s \equiv \mathbb{1}_{\{\neg\}}$ is a Nash equilibrium as required.

Step ($n > 0$): Assume w.l.o.g. that $t_0 \leq t_1 \leq \dots \leq t_{n-1}$. By the induction hypothesis, there exists a symmetric mixed-consumption Nash equilibrium s' in the $(n-1)$ -producers consumer game $(\mu|_{\cap [0, t_{n-1}]}; t_0, \dots, t_{n-2})$. If $\mu([t_{n-1}, 1]) = 0$, then we define a mixed-consumption profile s in $(\mu; \bar{t})$ s.t. $s|_{[0, t_{n-1}]} = s'|_{[0, t_{n-1}]}$, and $s|_{[t_{n-1}, 1]} \equiv \mathbb{1}_{\{n-1\}}$. As s' is symmetric in $(\mu|_{\cap [0, t_{n-1}]}; t_0, \dots, t_{n-2})$, so is s in $(\mu; \bar{t})$. As $\ell_j^s = \ell_j^{s'}$ for every $j \in \mathbb{P}_{n-1}$, by s' being a Nash equilibrium in $(\mu|_{\cap [0, t_{n-1}]}; t_0, \dots, t_{n-2})$ we have that no player of any type $d \in [0, t_{n-1})$ has any incentive to unilaterally deviate from s . As $\mu([t_{n-1}, 1]) = 0$, we have $\ell_{n-1}^s = 0$, and so players of types $d \in [t_{n-1}, 1]$ have no incentive to deviate from s either. Therefore, s is a symmetric Nash equilibrium as required, and the proof for this case is complete. Assume therefore, henceforth, that $\mu([t_{n-1}, 1]) > 0$.

Recall that $\ell_0^{s'}, \ell_1^{s'}, \dots, \ell_{n-2}^{s'}$ are the loads on producers in s' , and by slight abuse of notation, define $\ell_{n-1}^{s'} \triangleq 0 \leq \ell_{n-2}^{s'}$; by Lemma A.1, $\ell_0^{s'} \geq \ell_1^{s'} \geq \dots \geq \ell_{n-2}^{s'} \geq \ell_{n-1}^{s'}$. Let p be as in Lemma A.2 for $\ell_j = \ell_j^{s'}$ for every $j \in \mathbb{P}_n$, and for $m = \mu([t_{n-1}, 1]) > 0$. We define a mixed-consumption profile s in $(\mu; \bar{t})$ s.t. $s|_{[0, t_{n-1}]} = s'|_{[0, t_{n-1}]}$, and $s|_{[t_{n-1}, 1]} \equiv \frac{p}{\mu([t_{n-1}, 1])}$ (by Lemma A.2(1), indeed $\frac{p}{\mu([t_{n-1}, 1])} \in \Delta^{S_d}$ for all $d \geq t_{n-1}$). Once again, as

s' is symmetric, so is s . It remains to show that s is indeed a Nash equilibrium as required.

By definition of s' and of s , we have that $\ell_j^s = \ell_j^{s'} + p_j$ for every $j \in \mathbb{P}_n$. Let $d \in [0, t_{n-1})$. As $\ell_0^{s'}, \dots, \ell_{n-2}^{s'}$ and $\ell_0^s, \dots, \ell_{n-2}^s$ are both nonincreasing (the former by Lemma A.1, and the latter — by Lemma A.2(2)), and as S_d is the same in both $(\mu|_{\cap[0, t_{n-1}]}; t_0, \dots, t_{n-2})$ and $(\mu; \bar{t})$, we have that as no player of type d has any incentive to unilaterally deviate from s' in the former, neither does it from s in the latter. Let now $d \in [t_{n-1}, 1]$. For every $k \in \text{supp}(s(d))$, we have by definition $p_k > 0$, and so, by Lemma A.2(3), $\ell_k^s = \min_{j \in \mathbb{P}_n} \ell_j^s$, and the proof is complete.

The complexity claim follows as each inductive step requires $O(n)$ time — the time required to calculate p , by Lemma A.2. \square

COROLLARY A.4. *Let $h \in \mathbb{P}_n$, let s' be a mixed-consumption Nash equilibrium in the h -producers consumer game $(\mu|_{\cap[0, t_h]}; t_0, \dots, t_{h-1})$, and let s be the mixed-consumption Nash equilibrium in the n -producers consumer game $(\mu; \bar{t})$ constructed iteratively from s' as in the proof of Theorem 4.10. For every $0 \leq j < h$, we have $\ell_j^s \geq \ell_j^{s'}$, with equality if $\ell_{h-1}^s > \ell_h^s$.*

PROOF. By following the construction in the proof of Theorem 4.10, and by Lemma A.2(3). \square

Theorem 4.5 follows from Theorem 4.10 and from the following lemma.

LEMMA A.5 (THEOREM 4.10 \Rightarrow THEOREM 4.5). *If a mixed-consumption Nash equilibrium exists in the n -producers consumer game $(\mu; \bar{t})$, and if μ is atomless, then a pure-consumption Nash equilibrium exists in this game as well.*

PROOF. Assume w.l.o.g. that $t_0 \leq t_1 \leq \dots \leq t_{n-1}$. Let s be a mixed-consumption Nash equilibrium in the game $(\mu; \bar{t})$. For every $i \in \{0, \dots, n-2\}$, set $C^i \triangleq [t_i, t_{i+1})$, and set $C^{-1} \triangleq [0, t_0)$ and $C^{n-1} \triangleq [t_{n-1}, 1]$; note that $S_d = \{-\}$ for all $d \in C^{-1}$, and that $S_d = \{0, \dots, i\} \cup \{-\}$ for all $d \in C^i$, for every $i \in \mathbb{P}_n$. For every $i, j \in S$, define $p_j^i \triangleq \int_{C^i} s_j d\mu$; note that if $p_j^i > 0$, then $s_j(d) > 0$ for some $d \in C^i$. Let $i \in S$. We first consider the case in which either $\mu(C^i) > 0$ or $C^i = \emptyset$. In this case, as μ is atomless, there exists a partition of C^i into n pairwise-disjoint Lebesgue-measurable sets $(C_j^i)_{j \in S}$, s.t. $\mu(C_j^i) = p_j^i$ for all $j \in S$, and s.t. $C_j^i = \emptyset$ whenever $p_j^i = 0$. Otherwise, i.e. if $\mu(C^i) = 0$ yet $C^i \neq \emptyset$, then let $k \in S$ s.t. $s_k(d) > 0$ for some $d \in C^i$, and define $C_k^i \triangleq C^i$, and $C_j^i \triangleq \emptyset$ for every $j \in S \setminus \{k\}$. Note that in this case we also have that $(C_j^i)_{j \in S}$ is a partition of C^i and $\mu(C_j^i) = 0 = \int_{C^i} s_j d\mu = p_j^i$ for all $j \in S$.

We define a measurable function $s' : \mathcal{T} \rightarrow S$ by $s'|_{\cup_{i \in S} C_j^i} \equiv j$ for every $j \in S$. For every $j \in S$, we note that $\ell_j^{s'} = \mu(\cup_{i \in S} C_j^i) = \sum_{i \in S} p_j^i = \sum_{i \in S} \int_{C^i} s_j d\mu = \int_{\mathcal{T}} s_j d\mu = \ell_j^s$.

We conclude by showing that s' is indeed a pure-strategy profile, and moreover — a Nash equilibrium. Let $d \in \mathcal{T}$; by definition there exists $i \in S$ s.t. $d \in C_{s'(d)}^i \subseteq C^i$. As $C_{s'(d)}^i \neq \emptyset$, by definition of $C_{s'(d)}^i$ we have that $s_{s'(d)}(d') > 0$ for some $d' \in C^i$, and so $s'(d) \in S_d$. As by definition of C^i we have $S_d = S_{d'}$, we obtain $s'(d) \in S_d$, and so s' is a pure-strategy profile. Furthermore, as $s_{s'(d)}(d') > 0$, we obtain $\ell_{s'(d)}^{s'} = \ell_{s'(d)}^s = \min_{j \in S_{d'}} \ell_j^s = \min_{j \in S_d} \ell_j^s = \min_{j \in S_d} \ell_j^{s'}$, and the proof is complete. \square

PROOF OF THEOREM 4.11. Assume w.l.o.g. that $t_0 \leq t_1 \leq \dots \leq t_{n-1}$. Let s, s' be two mixed-consumption Nash equilibria in the game $(\mu; \bar{t})$. By definition of Nash equi-

librium, we have $s'_\neg = \mathbb{1}_{[0, t_0)} = s_\neg$, and so $\sum_{j=0}^{n-1} \ell_j^{s'} = \mu([t_0, 1]) = \sum_{j=0}^{n-1} \ell_j^s$. Assume for contradiction that there exists $j \in \mathbb{P}_n$ s.t. $\ell_j^{s'} \neq \ell_j^s$; let j be minimal with this property, and assume w.l.o.g. that $\ell_j^{s'} > \ell_j^s$.

Let $j \leq k < n$ be maximal s.t. $\ell_k^{s'} = \ell_j^{s'}$. By Lemma A.1, for every $j \leq i \leq k$, we have $\ell_i^{s'} = \ell_j^{s'} > \ell_j^s \geq \ell_i^s$. Therefore, and as $\ell_i^{s'} = \ell_i^s$ for every $0 \leq i < j$, we have $\sum_{j=0}^k \ell_j^{s'} > \sum_{j=0}^k \ell_j^s$. We thus obtain both that $k < n-1$, and that $\sum_{j=0}^k \ell_j^{s'} > \sum_{j=0}^k \ell_j^s \geq \mu([t_0, t_{k+1}])$. Therefore, $\sum_{j=k+1}^{n-1} \ell_j^{s'} < \mu([t_{k+1}, 1])$, and hence there exists $d \in [t_{k+1}, 1]$ s.t. $s'_i(d) > 0$ for some $0 \leq i \leq k$. As by Lemma A.1 we have $\ell_i^{s'} \geq \ell_k^{s'} > \ell_{k+1}^{s'}$, we conclude that s' is not a Nash equilibrium — a contradiction. \square

PROOF OF COROLLARY 4.12. Assume w.l.o.g. that $t_0 \leq t_1 \leq \dots \leq t_{n-1}$. Let s be a mixed-consumption Nash equilibrium in the game $(\mu; \bar{t})$, let $d \in \mathcal{T}$ and let $k \in \text{supp}(s(d))$. If $d < t_j$ for all $j \in \mathbb{P}_n$, then $k = \neg$ and so $\ell_k^s = \mu([0, t_0])$. Otherwise, $k \neq \neg$ and so $d \geq t_0$; let $i \in \mathbb{P}_n$ be largest s.t. $t_i \leq d$. By definition of s and by Lemma A.1, we have $\ell_k^s = \min\{\ell_j^s \mid t_j \leq d\} = \ell_i^s$. Either way (and by Theorem 4.11 when $k \neq \neg$), ℓ_k^s does not depend on the choice of s , as required. \square

From Algorithm 1, we obtain the following recursive identity for $\ell_k(\bar{t})$.

COROLLARY A.6. *If $t_0 \leq t_1 \leq \dots \leq t_{n-1}$, then defining $t_n \triangleq 2$, we have*

$$\ell_k(t_0, \dots, t_{n-1}) = \text{Max}_{k < j \leq n} \frac{\mu([t_0, t_j]) - \sum_{i=0}^{k-1} \ell_i(\bar{t})}{j - h} = \text{Max}_{k < j \leq n} \frac{\mu([0, t_j]) - \sum_{i \in \mathbb{P}_k \cup \{\neg\}} \ell_i(\bar{t})}{j - h}$$

(where by slight abuse of notation, μ is treated as a measure on $[0, 2]$ with support \mathcal{T}) for every $k \in \mathbb{P}_n$.

PROOF. A direct corollary of Algorithm 1, by considering two cases: in the first, either $k = 0$ or $\ell_k(\bar{t}) < \ell_{k-1}(\bar{t})$ (and so the given value k is the value of the variable k in some iteration of Algorithm 1); in the second, $k > 0$ and $\ell_k(\bar{t}) = \ell_{k-1}(\bar{t})$ (and so Algorithm 1 calculates both $\ell_k(\bar{t})$ and $\ell_{k-1}(\bar{t})$ in the same iteration of the while loop, and therefore they are identical; it is straightforward to verify that the expression in the statement evaluates to the same value for both $k-1$ and k in this case). \square

A.2. Analysis of ℓ

Before moving on to prove the results presented in Section 5, we now formalize three analytic properties of the function ℓ (defined in Definition 4.13), which we later utilize when proving the results of Section 5. The first property is that the load on a producer does not decrease if the producer raises the offered quality of service (i.e. lowers its strategy).

LEMMA A.7 (ℓ_j IS NONINCREASING IN t_j). *For every $j \in \mathbb{P}_n$ and for every $\bar{t} \in \mathcal{T}^{\mathbb{P}_n}$ and $t'_j \in \mathcal{T}$, if $t_j < t'_j$, then $\ell_j(\bar{t}_{-j}, t'_j) \leq \ell_j(\bar{t})$.*

PROOF. Let $\bar{t} \in \mathcal{T}^{\mathbb{P}_n}$, $k \in \mathbb{P}_n$ and $t'_k \in (t_k, 1]$. Assume w.l.o.g. that $t_0 \leq t_1 \leq \dots \leq t_{n-1} \in \mathcal{T}$. If $k \neq n-1$, then it is enough to consider the case $t_k < t'_k \leq t_{k+1}$. Let s be a mixed-consumption Nash equilibrium in the induced consumer game $(\mu; \bar{t})$. For every $j \in \mathbb{P}_n \setminus \{k\}$, define $t'_j \triangleq t_j$. Let s' be a mixed-consumption Nash equilibrium in $(\mu; \bar{t}')$, and assume for contradiction that $\ell_k^{s'} > \ell_k^s$. Let $i \in \mathbb{P}_n$ be maximal s.t. $\ell_i^{s'} = \ell_k^{s'}$; by definition, $i \geq k$.

We claim that $\ell_j^{s'} \geq \ell_j^s$ for every $0 \leq j \leq i$. Let $0 \leq h \leq i$ be minimal s.t. $\ell_h^s \leq \ell_k^{s'}$ (such h exists, and $h \leq k$, as $\ell_k^s < \ell_k^{s'}$); we will show that $\ell_j^{s'} \geq \ell_j^s$ separately for every $0 \leq j < h$ (if $h > 0$) and for every $h \leq j \leq i$. For every $h \leq j \leq i$, by Lemma A.1, by definition of i and by definition of h , we have $\ell_j^{s'} \geq \ell_i^{s'} = \ell_k^{s'} \geq \ell_h^s \geq \ell_j^s$, as required. We move on to show that $\ell_j^{s'} \geq \ell_j^s$ for every $0 \leq j < h$; assume that $h > 0$ (otherwise, there is nothing to show). Let $\tilde{\ell}_0, \dots, \tilde{\ell}_{h-1}$ be the loads on producers $0, \dots, h-1$ in a Nash equilibrium in the game $(\mu|_{\cap[0, t_h)}; t_0, \dots, t_{h-1})$; similarly, let $\tilde{\ell}'_0, \dots, \tilde{\ell}'_{h-1}$ be the loads on producers $0, \dots, h-1$ in a Nash equilibrium in the game $(\mu|_{\cap[0, t'_h)}; t'_0, \dots, t'_{h-1})$. As $h > 0$, by definition of h we have $\ell_{h-1}^s > \ell_h^s$; therefore, by Corollary A.4 and Theorem 4.11, we have that $\ell_j^s = \tilde{\ell}_j$ for every $0 \leq j < h$. By Corollary A.4 and Theorem 4.11, we obtain also that $\ell_j^{s'} \geq \tilde{\ell}'_j$ for every $0 \leq j < h$. As $k \geq h$, we have that $t'_j = t_j$ for every $0 \leq j < h$ and that $t'_{h-1} \geq t_{h-1}$; therefore, $\tilde{\ell}'_j \geq \tilde{\ell}_j$ for every $0 \leq j < h$. (This follows by tracing the construction in the proof of Theorem 4.10, as all inductive steps but the last are identical, while the last, examining Lemma A.2, increases each load by no less when computing $\tilde{\ell}'$ than when computing $\tilde{\ell}$; in the context of Fig. 1, pouring a additional nonnegative amount of liquid into the rightmost vessel does not cause the liquid level in any vessel to fall. Alternatively, this can also be seen by tracing Algorithm 1, as each iteration when computing $\tilde{\ell}'$ either computes the same load values for the producers as the corresponding iteration when computing $\tilde{\ell}$, or is the last, thus computing loads that are not lower than those computed for $\tilde{\ell}$.) Combining all of these, we have $\ell_j^{s'} \geq \tilde{\ell}'_j \geq \tilde{\ell}_j = \ell_j^s$ for every $0 \leq j < h$, as required.

We conclude that $\sum_{j=0}^i \ell_j^{s'} > \sum_{j=0}^i \ell_j^s$, as $\ell_j^{s'} \geq \ell_j^s$ for every $0 \leq j \leq i$, with a strict inequality for $j = k$. If $i = n-1$, then $\sum_{j=0}^{n-1} \ell_j^{s'} > \sum_{j=0}^{n-1} \ell_j^s = \mu([t_0, 1]) \geq \mu([t'_0, 1])$ — a contradiction; assume, therefore, that $i < n-1$. Hence, $\sum_{j=0}^i \ell_j^{s'} > \sum_{j=0}^i \ell_j^s \geq \mu([t_0, t_{i+1}]) \geq \mu([t'_0, t_{i+1}])$. Therefore, there exists $d \geq t'_{i+1}$ s.t. $s'_j(d) > 0$ for some $0 \leq j \leq i$, but by Lemma A.1 and by definition of i we notice that $\ell_j^{s'} \geq \ell_i^{s'} > \ell_{i+1}^{s'}$, so s' is not a Nash equilibrium in $(\mu; t')$ — a contradiction as well.

We note that an alternative proof may also be given via Algorithm 1 and Corollary A.6. \square

The second property is that the load on producer j cannot increase as a result of other producers moving closer to j 's quality of service.

LEMMA A.8 (ℓ_j IS WEAKLY QUASICONVEX IN t_k). *For every $j \in \mathbb{P}_n \cup \{-\}$ and $j \neq k \in \mathbb{P}_n$ and for every $\bar{t} \in \mathcal{T}^{\mathbb{P}_n}$ and $t'_k \in \mathcal{T}$, if $j \neq -$ and either $t_k < t'_k \leq t_j$ or $t_j \leq t'_k < t_k$, or if $j = -$ and $t'_k < t_k$, then $\ell_j(\bar{t}_{-k}, t'_k) \leq \ell_j(\bar{t})$.*

PROOF. Assume w.l.o.g. that $t_0 \leq t_1 \leq \dots \leq t_{n-1}$. If $k \neq n-1$, then it is enough to consider the case $t_k < t'_k \leq t_{k+1}$. Let s and s' be mixed-consumption Nash equilibria in $(\mu; \bar{t})$ and $(\mu; \bar{t}_{-k}, t'_k)$, respectively.

Assume for contradiction that $\ell_i^{s'} > \ell_i^s$ for some $k < i < n$, and let i be minimal with this property. Therefore (and by Lemma A.7 if $i = k+1$), $\ell_{i-1}^{s'} \leq \ell_{i-1}^s$. By Lemma A.1, we obtain $\ell_{i-1}^s \geq \ell_{i-1}^{s'} \geq \ell_i^{s'} > \ell_i^s$. Therefore, $s_j|_{[t_i, 1]} \equiv 0$ for every $0 \leq j < i$, and so $\sum_{j=i}^{n-1} \ell_j^{s'} \leq \mu([t_i, 1]) = \sum_{j=i}^{n-1} \ell_j^s$. Hence, and as $\ell_i^{s'} > \ell_i^s$, there exists $i < h < n$ s.t. $\ell_h^{s'} < \ell_h^s$ — let h be minimal with this property. Therefore, $\ell_{h-1}^{s'} \geq \ell_{h-1}^s$, and by Lemma A.1, we obtain $\ell_{h-1}^{s'} \geq \ell_{h-1}^s \geq \ell_h^s > \ell_h^{s'}$. By definition of h , we have that $\ell_j^{s'} \geq \ell_j^s$

for every $i \leq j < h$, with a strict inequality for $j = i$ by definition of i , and so $\sum_{j=i}^{h-1} \ell_j^{s'} > \sum_{j=i}^{h-1} \ell_j^s \geq \mu([t_i, t_h])$, with the last inequality since $s_j|_{[t_i, 1]} \equiv 0$ for every $0 \leq j < i$. Therefore, there exists $d \geq t_h$ s.t. $s'_j(d) > 0$ for some $i \leq j < h$, but by Lemma A.1, $\ell_j^{s'} \geq \ell_{h-1}^{s'} > \ell_h^s$, so s' is not a Nash equilibrium in $(\mu; \bar{t}_{-k}, t'_k)$ — a contradiction.

Assume now for contradiction that $\ell_i^{s'} < \ell_i^s$ for some $0 \leq i < k$, and let i be minimal with this property. As $k > 0$, we have $\ell_{-}^s = \mu([0, t_0]) = \ell_{-}^{s'}$. Therefore, $\sum_{j=0}^{n-1} \ell_j^{s'} = \sum_{j=0}^{n-1} \ell_j^s$ and by definition of i there exists $h \in \mathbb{P}_n$ s.t. $\ell_h^{s'} > \ell_h^s$ — let h be minimal with this property. By Lemma A.7 and by the first part of this proof, $h < k$. We now consider two cases: $h < i$ and $i < h$. We start with the case $h < i$. In this case, by definition of i and by Lemma A.1, $\ell_{i-1}^{s'} \geq \ell_{i-1}^s \geq \ell_i^s > \ell_i^{s'}$, and so, by Corollary A.4 and as $i < k$, $\ell_h^{s'} = \ell_h(\mu|_{\cap[0, t_i]}; t_0, \dots, t_{i-1}) \leq \ell_h^s$ — a contradiction. Similarly, if $i < h$, then by definition of h and by Lemma A.1, $\ell_{h-1}^s \geq \ell_{h-1}^{s'} \geq \ell_h^{s'} > \ell_h^s$, and so, by Corollary A.4 and as $h < k$, $\ell_i^s = \ell_i(\mu|_{\cap[0, t_h]}; t_0, \dots, t_{h-1}) \leq \ell_i^{s'}$ — a contradiction as well.

We conclude by examining the effect on ℓ_{-} . If $k \neq 0$, then let $t'_0 \triangleq t_0$. Regardless of the value of k , we have $t'_0 \geq t_0$. By definition of s, s' , we have $\ell_{-}^s = \mu([0, t_0]) \leq \mu([0, t'_0]) = \ell_{-}^{s'}$.

We note that an alternative proof may also be given via Algorithm 1 and Corollary A.6. \square

Finally, the last property is that small perturbations in the producers' strategies result in quantifiably-small changes in the loads on the various producers.

LEMMA A.9 (*ℓ IS LIPSCHITZ IN EACH COORDINATE WITH LIPSCHITZ CONSTANT 1*). *For every $j, k \in \mathbb{P}_n$ and for every $\bar{t} \in \mathcal{T}^{\mathbb{P}_n}$ and $t'_k \in \mathcal{T}$, if $t_k < t'_k$, then $|\ell_j(\bar{t}_{-k}, t'_k) - \ell_j(\bar{t})| \leq \mu([t_k, t'_k])$.*

PROOF. Assume w.l.o.g. that $t_0 \leq t_1 \leq \dots \leq t_{n-1}$. If $k \neq n-1$, then it is enough to consider the case $t_k < t'_k \leq t_{k+1}$. Let s and s' be mixed-consumption Nash equilibria in $(\mu; \bar{t})$ and $(\mu; \bar{t}_{-k}, t'_k)$, respectively. Define $S_k \triangleq \mathbb{P}_k \cup \{-\}$. We start by showing that $\sum_{j \in S_k} \ell_j^{s'} - \sum_{j \in S_k} \ell_j^s \leq \mu([t_k, t'_k])$.

If $k = 0$, then this claim holds, as in this case $S_k = \{-\}$, and by definition of s and s' we have $\ell_{-}^{s'} - \ell_{-}^s = \mu([0, t'_k]) - \mu([0, t_k]) = \mu([t_k, t'_k])$. Assume therefore that $k > 0$ and assume for contradiction that $\sum_{j \in S_k} \ell_j^{s'} - \sum_{j \in S_k} \ell_j^s > \mu([t_k, t'_k])$. As $k > 0$, we have $\ell_{-}^{s'} = \mu([0, t_0]) = \ell_{-}^s$, and so $\sum_{j=0}^{k-1} \ell_j^{s'} - \sum_{j=0}^{k-1} \ell_j^s > \mu([t_k, t'_k])$ as well. Let $i \in \mathbb{P}_n$ be maximal s.t. $\ell_i^{s'} = \ell_{k-1}^{s'}$. By definition, $i \geq k-1$. For every $k \leq j \leq i$, by Lemmas A.1 and A.8, we have $\ell_j^{s'} \geq \ell_i^{s'} = \ell_{k-1}^{s'} \geq \ell_{k-1}^s \geq \ell_j^s$. Therefore, we have $\sum_{j=0}^i \ell_j^{s'} - \sum_{j=0}^i \ell_j^s > \mu([t_k, t'_k])$. If $i = n-1$, then we obtain $\sum_{j=0}^{n-1} \ell_j^{s'} > \sum_{j=0}^{n-1} \ell_j^s = \mu([t_0, 1])$ — a contradiction; assume, therefore, that $i < n-1$. If $i+1 \neq k$, then let $t'_{i+1} \triangleq t_{i+1}$. Hence,

$$\begin{aligned} \sum_{j=0}^i \ell_j^{s'} &> \sum_{j=0}^i \ell_j^s + \mu([t_k, t'_k]) \geq \mu([0, t_{i+1}]) + \mu([t_k, t'_k]) = \\ &= \begin{cases} \mu([0, t_k]) + \mu([t_k, t'_k]) = \mu([0, t'_{i+1}]) & i+1 = k \\ \mu([0, t'_{i+1}]) + \mu([t_k, t'_k]) \geq \mu([0, t'_{i+1}]) & i+1 \neq k \end{cases} \end{aligned}$$

Thus, there exists $d \geq t'_{i+1}$ s.t. $s'_j(d) > 0$ for some $0 \leq j \leq i$, but by Lemma A.1 and

by definition of i we notice that $\ell_j^{s'} \geq \ell_i^{s'} > \ell_{i+1}^{s'}$, so s' is not a Nash equilibrium in $(\mu; \bar{t}_{-k}, t'_k)$ — a contradiction as well.

As $\sum_{j \in S_k} \ell_j^{s'} - \sum_{j \in S_k} \ell_j^s \leq \mu([t_k, t'_k])$ and as by Lemma A.8 $\ell_j^{s'} \geq \ell_j^s$ for every $j \in S_k$, we obtain $0 \leq \ell_j^{s'} - \ell_j^s \leq \mu([t_k, t'_k])$ for every $j \in S_k$. Similarly, As $\sum_{j \in S_k} \ell_j^{s'} - \sum_{j \in S_k} \ell_j^s \leq \mu([t_k, t'_k])$ and as $\sum_{j \in S} \ell_j^{s'} = \mu(\mathcal{T}) = \sum_{j \in S} \ell_j^s$, we have $\sum_{j=k}^{n-1} \ell_j^s - \sum_{j=k}^{n-1} \ell_j^{s'} \leq \mu([t_k, t'_k])$, and as by Lemmas A.7 and A.8 $\ell_j^{s'} \leq \ell_j^s$ for every $k \leq j < n$, we obtain $0 \leq \ell_j^s - \ell_j^{s'} \leq \mu([t_k, t'_k])$ for every $k \leq j < n$, and the proof is complete.

We note that an alternative proof may also be given via Algorithm 1 / Corollary A.6. \square

A.3. Proofs and Additional Results for Section 5.1

A.3.1. Proofs and Additional Results for Section 5.1.1

LEMMA A.10. *Let $j \in \mathbb{P}_n$ and $t_j \in \mathcal{T}$. For every $\bar{t}_{-j} \in \mathcal{T}^{\mathbb{P}_n \setminus \{j\}}$, we have $\ell_j(\bar{t}) \geq \frac{\mu([t_j, 1])}{n}$, which constitutes a tight bound.*

PROOF. Let s be mixed-consumption Nash equilibrium in $(\mu; \bar{t})$. If $\mu([t_j, 1]) = 0$, then $\ell_j^s = 0$ and the claim trivially holds. Assume therefore that $\mu([t_j, 1]) > 0$. Let $k \in \arg \max_{i \in \mathbb{P}_n} \int_{[t_j, 1]} s_i d\mu$. By definition of s , $\int_{[t_j, 1]} s_{-} d\mu = 0$, and so $\int_{[t_j, 1]} s_k d\mu \geq \frac{\mu([t_j, 1])}{n}$ by definition of k . Since $\int_{[t_j, 1]} s_k d\mu \geq \frac{\mu([t_j, 1])}{n} > 0$, there exists $d \geq t_j$ s.t. $s_k(d) > 0$. As $j \in S_d$, by definition of Nash equilibrium we have $\int_{[t_j, 1]} s_j d\mu = \ell_j^s \geq \ell_k^s \geq \int_{[t_j, 1]} s_k d\mu \geq \frac{\mu([t_j, 1])}{n}$.

Alternatively, by Lemma A.8, $\ell_j(\bar{t})$ is minimal given t_j when $t_i = t_j$ for every $i \in \mathbb{P}_n \setminus \{j\}$. By Theorem 4.11, by anonymity, and by definition of Nash equilibrium, the load on each producer in this case is exactly $\frac{\mu([t_j, 1])}{n}$. \square

COROLLARY A.11. *Let $\bar{t} \in \mathcal{T}^{\mathbb{P}_n}$. For every $j \in \mathbb{P}_n$, if $t_j = 0$, then $\ell_j(\bar{t}) \geq \frac{\mu(\mathcal{T})}{n}$.*

PROOF. A direct corollary of Lemma A.10. \square

Definition A.12 (Domination). Let t, t' be strategies in (n, μ, \succeq_C) .

- We say that t *weakly dominates* t' if t is a safe alternative to t' and moreover, there exists some strategy profile for all but one of the producers, s.t. playing t gives the remaining producer strictly higher utility than playing t' .
- We say that t *strongly dominates* t' if for every strategy profile for all but one of the producers, playing t gives the remaining producer strictly higher utility than playing t' .

LEMMA A.13 (DOMINATION). *Let $t < t' \in \mathcal{T}$ be strategies in (n, μ, \succeq_C) .*

- (1) t is a safe alternative to t' .
- (2) t weakly dominates t' iff $\mu([t, t']) > 0$.
- (3) t strongly dominates t' iff $\mu([t', 1]) < \frac{\mu([t, 1])}{n}$.

PROOF. Part 1 follows from Lemma A.7.

We move on to proving Part 2. \Rightarrow : Assume that $\mu([t, t']) = 0$; by Algorithm 1 / Corollary A.6, t and t' are equivalent, and *a fortiori* t' is a safe alternative to t . Nonetheless, we now also directly show that t and t' are equivalent. Let $\bar{t}_{-0} \in \mathcal{T}^{\mathbb{P}_n \setminus \{0\}}$, and let s be a mixed-consumption Nash equilibrium in $(\mu; \bar{t}_{-0}, t)$. Let s' be the mixed-consumption profile in $(\mu; \bar{t}_{-0}, t')$ s.t. $s'|_{\mathcal{T} \setminus [t, t']} = s|_{\mathcal{T} \setminus [t, t']}$ and s.t. for every $d \in [t, t']$, if $S_d = \{-\}$ w.r.t.

$(\mu; \bar{t}_{-0}, t')$ then $s'(d) = \mathbb{1}_{\{-\}}$, and otherwise $s'(d) = \mathbb{1}_{\{j\}}$ for some $j \in \arg \text{Max}_{j \in S_d} t_j$. As $s = s'$ almost everywhere w.r.t. μ , we have that $\ell_j^{s'} = \ell_j^s$ for all $j \in \mathbb{P}_n$. By definition of s , therefore no type $d \in \mathcal{T} \setminus [t, t']$ has any incentive to deviate from s' in the latter game. By Lemma A.1 (for s) and by definition of s' , neither does any type $d \in [t, t']$ have any incentive to deviate from s' in the latter game. Therefore, s' is a Nash equilibrium in $(\mu; \bar{t}_{-0}, t')$. As in particular $\ell_0^{s'} = \ell_0^s$, the proof of this direction is complete.

\Leftarrow : Assume that $\mu([t, t']) > 0$; we will show that t' is not a safe alternative to t . Define $a \triangleq \mu([t, t']) > 0$, $b \triangleq \mu([t', 1])$ and $c \triangleq \mu(\{1\})$. By Algorithm 1, we have $\ell_0(t', 1, 1, \dots, 1) = \text{Max}\{b, \frac{b+c}{n}\} < \text{Max}\{a+b, \frac{a+b+c}{n}\} = \ell_0(t, 1, 1, \dots, 1)$, and the proof of this direction is complete as well.

We conclude by proving Part 3. \Rightarrow : Assume that $\mu([t', 1]) \geq \frac{\mu([t, 1])}{n}$. Therefore, $\mu([t, t']) \leq \frac{n-1}{n} \cdot \mu([t, 1])$, and hence $\frac{\mu([t, t'])}{n-1} \leq \frac{\mu([t, 1])}{n}$. By Theorem 4.11, by anonymity, and by definition of Nash equilibrium, the load on every producer, and in particular on producer 0, in a Nash equilibrium in the n -producer game $(\mu; t, t, \dots, t)$ is $\frac{\mu([t, 1])}{n}$. As $\text{Max}\{\frac{\mu([t, t'])}{n-1}, \frac{\mu([t, 1])}{n}\} = \frac{\mu([t, 1])}{n}$, by Algorithm 1 the load on every producer, and in particular on producer 0, in a Nash equilibrium in the game $(\mu|_{\cap[0, t']}; t', t, t, \dots, t)$ is $\frac{\mu([t, 1])}{n}$ as well, as required.

\Leftarrow : Assume that $\mu([t', 1]) < \frac{\mu([t, 1])}{n}$. Let $\bar{t}_{-0} \in \mathcal{T}^{\mathbb{P}_n} \setminus \{0\}$. By Lemma A.10 and by definition of legal strategies in the consumer game, we obtain $\ell_0(\bar{t}_{-0}, t) \geq \frac{\mu([t, 1])}{n} > \mu([t', 1]) \geq \ell_0(\bar{t}_{-0}, t')$. \square

PROOF OF THEOREM 5.3. The first statement is a direct corollary of Lemma A.13, and the second — of Corollary A.11. \square

LEMMA A.14. *Let $\bar{t} \in \mathcal{T}^{\mathbb{P}_n}$ and let s be a mixed-consumption profile in the consumer game $(\mu; \bar{t})$. If $\ell_j^s = \frac{\mu(\mathcal{T})}{n}$ for every $j \in \mathbb{P}_n$, and if $s_-(d) = 0$ whenever $S_d \neq \{-\}$, then s constitutes a Nash equilibrium in this game.*

PROOF. Directly from definition of mixed-consumption Nash equilibrium, no consumer has any incentive to unilaterally deviate from s . \square

PROOF OF THEOREM 5.4. The first part (\Rightarrow) follows directly from Corollary A.11. For the second part (\Leftarrow), let \bar{t} be a pure-strategy profile in (n, μ, \succeq_C) s.t. $\ell_j(\bar{t}) = \frac{\mu(\mathcal{T})}{n}$ for every $j \in \mathbb{P}_n$. Assume w.l.o.g. that $t_0 \leq \dots \leq t_{n-1}$, and let s be a mixed-consumption Nash equilibrium in the induced consumer game $(\mu; \bar{t})$; therefore, $\ell_j^s = \ell_j(\bar{t}) = \frac{\mu(\mathcal{T})}{n}$ for every $j \in \mathbb{P}_n$. Let $k \in \mathbb{P}_n$ and let $t'_k \in \mathcal{T}$; we will show that producer k has no incentive to deviate to t'_k from t_k . If $t_k < t'_k$, then this follows directly from Lemma A.13(1). We therefore consider the case in which $t'_k < t_k$. Let $N \triangleq [t'_k, t_0]$ (if $t_0 \leq t'_k$, then $N = \emptyset$). By definition of s , we have that $\mu(N) \leq \mu([0, t_0]) = \ell_0^s = 0$. We define a mixed-consumption profile s' in the consumer game $(\mu; \bar{t}_{-k}, t'_k)$ by $s'|_N \equiv \mathbb{1}_{\{k\}}$ and $s'|_{\mathcal{T} \setminus N} = s|_{\mathcal{T} \setminus N}$ (if $N = \emptyset$, then $s' = s$). (This indeed is a mixed-consumption profile since $t'_k < t_k$ and by definition of N .) As $\mu(N) = 0$, we have $\ell_j^{s'} = \ell_j^s = \frac{\mu(\mathcal{T})}{n}$ for every $j \in \mathbb{P}_n$. Thus, by Lemma A.14 and by definition of s' via N , we conclude that s is a Nash equilibrium in $(\mu; \bar{t}_{-k}, t'_k)$, and so k has no incentive to deviate to t'_k from t_k in this case either. \square

COROLLARY A.15 (LEAST/MOST NASH EQUILIBRIUM LOAD). *Let $t_0 \leq \dots \leq t_{n-1} \in \mathcal{T}$. The pure-strategy profile \bar{t} constitutes a Nash equilibrium in (n, μ, \succeq_C) iff either of the following equivalent conditions hold.*

- $\ell_{n-1}(\bar{t}) = \frac{\mu(\mathcal{T})}{n}$.
- $\mu([0, t_0]) = 0$ and $\ell_0(\bar{t}) = \frac{\mu(\mathcal{T})}{n}$.

PROOF. A direct corollary of Theorem 5.4 and Lemma A.1. \square

PROOF OF THEOREM 5.5. \Rightarrow : Assume that \bar{t} constitutes a pure-strategy Nash equilibrium in (n, μ, \succeq_C) . Let s be a mixed-consumption Nash equilibrium in the induced consumer game $(\mu; \bar{t})$. For every $j \in \mathbb{P}_n$, by Theorem 5.4, we obtain $\mu([t_j, 1]) \geq \sum_{k=j}^{n-1} \ell_k^s = \frac{n-j}{n} \cdot \mu(\mathcal{T})$, and so $\mu([0, t_j]) \leq \frac{j}{n} \cdot \mu(\mathcal{T})$.

\Leftarrow : Assume that $\mu([0, t_j]) \leq \frac{j}{n} \cdot \mu(\mathcal{T})$ for every $j \in \mathbb{P}_n$; in particular, $\mu([0, t_0]) = 0$. Let s be a mixed-consumption profile in the induced consumer game $(\mu; \bar{t})$. By Theorem 5.4, it is enough to show that $\ell_j^s = \frac{\mu(\mathcal{T})}{n}$ for every $j \in \mathbb{P}_n$. Assume for contradiction that this is not the case. Therefore, as $\ell_{-}^s = \mu([0, t_0]) = 0$, there exists $k \in \mathbb{P}_n$ s.t. $\ell_k^s > \frac{\mu(\mathcal{T})}{n}$; let k be maximal with this property. By Lemma A.1, we have $\sum_{j=0}^k \ell_j^s \geq (k+1) \cdot \ell_k^s > \frac{(k+1) \cdot \mu(\mathcal{T})}{n} \geq \mu([0, t_{k+1}])$. Therefore, there exists $d \geq t_{k+1}$ s.t. $s'_j(d) > 0$ for some $0 \leq j \leq k$, but by definition of k we notice that $\ell_k^{s'} > \frac{\mu(\mathcal{T})}{n} \geq \ell_{k+1}^s$, so s' is not a Nash equilibrium — a contradiction.

We note that an alternative proof of the second direction (\Leftarrow) may also be given via Algorithm 1 / Corollary A.6. \square

The second direction (\Leftarrow) of Theorem 5.5 can also be proven constructively. Such a proof is quite tedious in the general case, but simplifies greatly when μ is atomless. E.g. if $\mu = U(\mathcal{T})$ the uniform measure, then whenever $t_0 \leq \dots \leq t_{n-1}$ meet the conditions of Theorem 5.5, then a Nash equilibrium can be formed by allocating the market as follows: every $d \in [0, \frac{1}{n})$ plays the pure strategy $0 \in \mathbb{P}_n$ (the conditions of Theorem 5.5 guarantee that this is a legal strategy for all such d), every $d \in [\frac{1}{n}, \frac{2}{n})$ plays the pure strategy $1 \in \mathbb{P}_n$ (once again, the conditions of Theorem 5.5 guarantee that this is a legal strategy for all such d), and so on, until every $d \in [\frac{n-1}{n}, 1]$, playing the pure strategy $n-1 \in \mathbb{P}_n$. (We remark that if $t_j = \frac{j}{n}$ for every $j \in \mathbb{P}_n$, then this is in fact the unique Nash equilibrium among consumers, up to modifications of measure zero. More about this allocation and Nash equilibrium — in Theorem 5.31 and Remark 5.32 in Appendix A.4.) The load on each producer in this case is precisely $\frac{1}{n}$, and by Lemma A.14 and Theorem 5.4, the proof is complete.

PROOF OF THEOREM 5.6. Assume w.l.o.g. that $t_0 \leq t_1 \leq \dots \leq t_{n-1}$ and assume for contradiction that $P \subseteq \mathbb{P}_n$ and $\bar{t}' \in \mathcal{T}^P$ as in the statement exist. Let s be a mixed-consumption Nash equilibrium in $(\mu; \bar{t}_{-P}, \bar{t}')$. As there exists $j \in P$ s.t. $\ell_j^s > \ell_j(\bar{t}) = \frac{\mu(\mathcal{T})}{n}$ (with the last equality by Theorem 5.4), and as $\sum_{j=0}^{n-1} \ell_j^s \leq \mu(\mathcal{T})$, there thus exists a producer $k \in \mathbb{P}_n$ s.t. $\ell_k^s < \frac{\mu(\mathcal{T})}{n}$ — let $\emptyset \neq K \subseteq \mathbb{P}_n$ be the set of all such producers; by definition of P and by Theorem 5.4, we have $K \subseteq \mathbb{P}_n \setminus P$. By Lemma A.1, $K = \{n-1, n-2, \dots, n-|K|\}$. Therefore, and by Theorem 5.5, $\sum_{k=n-|K|}^{n-1} \ell_k^s < \frac{|K| \cdot \mu(\mathcal{T})}{n} \leq \mu([t_{n-|K|}, 1])$, and so there exists $d \in [t_{n-|K|}, 1]$ s.t. $s_j(d) > 0$ for some $0 \leq j < n-|K|$, but by definition of K we notice that $\ell_j^{s'} \geq \frac{\mu(\mathcal{T})}{n} > \ell_{n-|K|}^{s'}$, and so (as $n-|K| \notin P$) we have that s is not a Nash equilibrium in $(\mu; \bar{t}_{-P}, \bar{t}')$ — a contradiction. \square

PROOF OF THEOREMS 5.7 AND 5.8. Theorem 5.7(1) follows directly from Theorem 5.3. We move on to proving Theorem 5.8 and Theorem 5.7(2). Let \bar{p} be a mixed-strategy Nash equilibrium in (n, μ, \succeq_C) . By Corollary A.11, $E[\ell_j(\bar{p})] \geq \frac{\mu(\mathcal{T})}{n}$ for

every $j \in \mathbb{P}_n$; since $\sum_{j=0}^{n-1} \ell_j(\bar{p}) \leq \mu(\mathcal{T})$, and by linearity of expectation, we obtain $E[\ell_j(\bar{p})] = \frac{\mu(\mathcal{T})}{n}$ for every $j \in \mathbb{P}_n$. Let $j \in \mathbb{P}_n$. By Lemma A.13(1), we have $\ell_j(\bar{p}_{-j}, \mathbb{1}_{\{0\}}) \geq \ell_j(\bar{p})$. As j has no incentive to deviate from p_j to playing $0 \in \mathcal{T}$, we thus have that $\ell_j(\bar{p}_{-j}, \mathbb{1}_{\{0\}}) = \ell_j(\bar{p})$ with probability 1. By Corollary A.11, we have that $\ell_j(\bar{p}_{-j}, \mathbb{1}_{\{0\}}) \geq \frac{\mu(\mathcal{T})}{n}$, and so $\ell_j(\bar{p}) \geq \frac{\mu(\mathcal{T})}{n}$ with probability 1. As $E[\ell_j(\bar{p})] = \frac{\mu(\mathcal{T})}{n}$, we have that $\ell_j(\bar{p}) = \frac{\mu(\mathcal{T})}{n}$ with probability 1.

Let now \bar{p} be a mixed-strategy profile in (n, μ, \succeq_C) , s.t. $\ell_j(\bar{p}) = \frac{\mu(\mathcal{T})}{n}$ with probability 1 for every $j \in \mathbb{P}_n$. By Theorem 5.4, the resulting realization is a pure-strategy Nash equilibrium with probability 1, and so with probability 1 no ex-post regret exists and *a fortiori* \bar{p} is a Nash equilibrium.

We move on to proving Theorem 5.7(3). Let \bar{p} be a mixed-strategy Nash equilibrium in (n, μ, \succeq_C) . We iteratively build a permutation $\pi \in \mathbb{P}_n!$ s.t. $\mu([0, \text{Max supp}(p_{\pi(j)}))) \leq \frac{j}{n} \cdot \mu(\mathcal{T})$ for every $j \in \mathbb{P}_n$. Let $k \in \mathbb{P}_n$ and assume that $\pi(j)$ has been defined for all $0 \leq j < k$. Define $U \triangleq \mathbb{P}_n \setminus \{\pi(0), \dots, \pi(k-1)\}$ and let t_{\min} be a random variable denoting the numerically-smallest strategy realization of U , i.e. $t_{\min} \triangleq \min_{j \in U} p_j$. By Theorem 5.7(2), with probability 1 we have $\sum_{j \in U} \ell_j(\bar{p}) = \frac{n-k}{n} \cdot \mu(\mathcal{T})$; therefore, with probability 1 we have $\mu([0, t_{\min}]) \leq \frac{k}{n} \cdot \mu(\mathcal{T})$. Hence, by independence, it is not possible that $P[\mu([0, p_j]) > \frac{k}{n} \cdot \mu(\mathcal{T})] > 0$ for every $j \in U$. Therefore, there exists $\pi(k) \in U$ s.t. $P[\mu([0, p_{\pi(k)}]) \leq \frac{k}{n} \cdot \mu(\mathcal{T})] = 1$, and so $\mu([0, \text{Max supp}(p_{\pi(k)}))) \leq \frac{k}{n} \cdot \mu(\mathcal{T})$ and the construction is complete.

Let now \bar{p} be a mixed-strategy profile in (n, μ, \succeq_C) , s.t. $\mu([0, \text{Max supp}(p_{\pi(j)}))) \leq \frac{j}{n} \cdot \mu(\mathcal{T})$ for every $j \in \mathbb{P}_n$. Therefore, $\mu([0, p_{\pi(j)}]) \leq \frac{j}{n} \cdot \mu(\mathcal{T})$ for every $j \in \mathbb{P}_n$ with probability 1. By Theorem 5.5, the resulting realization is a pure-strategy Nash equilibrium with probability 1, and so with probability 1 no ex-post regret exists and *a fortiori* \bar{p} is a Nash equilibrium.

We conclude by proving Theorem 5.7(4). Let \bar{p} be a mixed-strategy Nash equilibrium in (n, μ, \succeq_C) . By Theorems 5.6 and 5.8, a realization of \bar{p} is with probability 1 a super-strong equilibrium, and so *a fortiori* \bar{p} is a super-strong Nash equilibrium. \square

A.3.2. Proofs and Additional Results for Section 5.1.2

LEMMA A.16. *Let $\bar{t} \in \mathcal{T}^{\mathbb{P}_n}$ be a pure-strategy profile, let $h \in \arg \text{Max}_{j \in \mathbb{P}_n} t_j$ and let $k \in \mathbb{P}_n$. If \bar{t} is not a Nash equilibrium in (n, μ, \succeq_C) , but nonetheless t_k is a best response to \bar{t}_{-k} , then both $\ell_h(\bar{t}) < \ell_k(\bar{t})$ and $\mu([t_k, t_h]) \geq \frac{\mu(\mathcal{T})}{n}$.*

PROOF. As t_k is a best response to \bar{t}_{-k} and by Corollary A.11, $\ell_k(\bar{t}) \geq \frac{\mu(\mathcal{T})}{n}$. As \bar{t} is not a Nash equilibrium, by Corollary A.15 we have $\ell_h(\bar{t}) < \frac{\mu(\mathcal{T})}{n}$, and so $\ell_h(\bar{t}) < \frac{\mu(\mathcal{T})}{n} \leq \ell_k(\bar{t})$. Therefore, in any mixed-consumption Nash equilibrium in $(\mu; \bar{t})$, no consumer with type $d \geq t_h$ consumes a positive amount from producer k , and so $\mu([t_k, t_h]) \geq \ell_k(\bar{t}) \geq \frac{\mu(\mathcal{T})}{n}$, as required. \square

Definition A.17. Let $\bar{t} \in \mathcal{T}^{\mathbb{P}_n}$ be a pure-strategy profile in (n, μ, \succeq_C) .

— For every $q \in \{0, \dots, n\}$, we define

$$Q_q(\bar{t}) \triangleq \left\{ j \in \mathbb{P}_n \mid \mu([0, t_j]) \in \left(\frac{q-1}{n} \cdot \mu(\mathcal{T}), \frac{q}{n} \cdot \mu(\mathcal{T}) \right] \right\} \subseteq \mathbb{P}_n.$$

— We define $M(\bar{t}) \triangleq \text{Max}\{q \in \{0, \dots, n\} \mid Q_q(\bar{t}) \neq \emptyset\}$.

Remark A.18.

- $(Q_q(\bar{t}))_{0 \leq q \leq n}$ is a partition of \mathbb{P}_n .
- When the CDF of μ is continuous and strictly increasing, then $Q_0(\bar{t})$ is the set of producers with strategy 0, while $Q_q(\bar{t})$ for $0 < q \leq n$ is the set of producers whose strategies lie in the q^{th} $1/n$ of \mathcal{T} (as measured by μ), i.e. above the $(q-1)^{\text{th}}$ n -tile yet not above the q^{th} n -tile; for such a CDF, $M(\bar{t})$ is the index of the $1/n$ of \mathcal{T} containing the numerically-largest strategy (or 0 if all strategies are 0), i.e. it is the index of the lowest n -tile above which no strategies lie.

LEMMA A.19. *Let $\bar{t} \in \mathcal{T}^{\mathbb{P}_n}$ and let t'_k be a best response to \bar{t}_{-k} . If \bar{t} is not a Nash equilibrium, then $\mu([0, t'_k]) \leq \frac{M(\bar{t})-1}{n} \cdot \mu(\mathcal{T})$.*

PROOF. Let $h \in \arg \text{Max}_{j \in \mathbb{P}_n} t'_j$, where $\bar{t}'_{-k} \triangleq \bar{t}_{-k}$. We consider two cases. If (\bar{t}_{-k}, t'_k) is not a Nash equilibrium, then by Lemma A.16 we have $\mu([t'_k, t_h]) \geq \frac{\mu(\mathcal{T})}{n}$, and so $\mu([0, t'_k]) \leq \mu([0, t_h]) - \frac{\mu(\mathcal{T})}{n} \leq \frac{M(\bar{t})-1}{n} \cdot \mu(\mathcal{T})$, as required. Otherwise, (\bar{t}_{-k}, t'_k) is a Nash equilibrium while \bar{t} is not. Therefore, by Theorem 5.5, there exists $\tilde{j} \in \{0, \dots, n-1\}$ s.t. $\mu([0, t'_k]) \leq \frac{\tilde{j}}{n} \cdot \mu(\mathcal{T}) < \mu([0, t_k])$. As $\mu([0, t_k]) \leq \frac{M(\bar{t})}{n} \cdot \mu(\mathcal{T})$, we thus have $\tilde{j} < M(\bar{t})$, and so $\mu([0, t'_k]) \leq \frac{M(\bar{t})-1}{n} \cdot \mu(\mathcal{T})$, as required. \square

LEMMA A.20. *Let $\delta > 0$, let $(\bar{t}^i, P_i)_{i=0}^\infty$ be a δ -better-response dynamic in (n, μ, \succeq_C) and let $i \in \mathbb{N}$. If \bar{t}^i is not a Nash equilibrium, then all of the following hold.*

- (1) $M(\bar{t}^{i+1}) \leq M(\bar{t}^i)$.
- (2) $Q_{M(\bar{t}^i)}(\bar{t}^{i+1}) \subseteq Q_{M(\bar{t}^i)}(\bar{t}^i)$.
- (3) $\mu([t_j^{i+1}, t_j^i]) \geq \delta$, for every $j \in P_i \cap Q_{M(\bar{t}^i)}(\bar{t}^{i+1})$.

Remark A.21. Finer analysis of similar nature may be used to show that δ may be replaced with $(n+1 - M(\bar{t}^i)) \cdot \delta$ in Lemma A.20(3).

PROOF OF LEMMA A.20. We commence by proving Part 1. Let $h \in \arg \text{Max}_{j \in \mathbb{P}_n} t_j^i$. It is enough to show that $t_k^{i+1} < t_h^i$ for every $k \in P_i$. We consider two cases. If t_k^i is not a best response to \bar{t}_{-k}^i , then by definition of δ -better-response dynamics and by Lemma A.13(1), $t_k^{i+1} < t_k^i \leq t_h^i$. Otherwise, t_k^i , and hence also t_k^{i+1} , are best responses to \bar{t}_{-k}^i . Therefore, by Lemma A.16, $\ell_h(\bar{t}^i) < \ell_k(\bar{t}^i) = \ell_k(\bar{t}_{-k}^i, t_k^{i+1})$; in particular, $k \neq h$. By anonymity and by Lemma A.8 (for $j = h$), we have $\ell_k(\bar{t}_{-k}^i, t_h^i) = \ell_h(\bar{t}_{-k}^i, t_h^i) \leq \ell_h(\bar{t}^i) < \ell_k(\bar{t}_{-k}^i, t_k^{i+1})$. Therefore, by Lemma A.7, $t_k^{i+1} < t_h^i$ in this case as well, as required.

We now proceed to prove Part 2. Let $k \in \mathbb{P}_n \setminus Q_{M(\bar{t}^i)}(\bar{t}^i)$; we must show that $k \notin Q_{M(\bar{t}^i)}(\bar{t}^{i+1})$. It is enough to consider the scenario in which $k \in P_i$, and to show that under this condition, $\mu([0, t_k^{i+1}]) \leq \frac{M(\bar{t}^i)-1}{n} \cdot \mu(\mathcal{T})$. Once again, we consider two cases. If t_k^i is not a best response to \bar{t}_{-k}^i , then by definition of δ -better-response dynamics and by Lemma A.13(1), $t_k^{i+1} < t_k^i$ and so $\mu([0, t_k^{i+1}]) \leq \mu([0, t_k^i]) \leq \frac{M(\bar{t}^i)-1}{n} \cdot \mu(\mathcal{T})$, as required. Otherwise, t_k^i , and hence also t_k^{i+1} , are best responses to \bar{t}_{-k}^i . By Lemma A.19, in this case we have $\mu([0, t_k^{i+1}]) \leq \frac{M(\bar{t}^i)-1}{n} \cdot \mu(\mathcal{T})$ as well, as required.

We conclude by proving Part 3. Let $k \in P_i \cap Q_{M(\bar{t}^i)}(\bar{t}^{i+1})$. As $\mu([0, t_k^{i+1}]) > \frac{M(\bar{t}^i)-1}{n} \cdot \mu(\mathcal{T})$, by Lemma A.19 we have that t_k^{i+1} is not a best response to \bar{t}_{-k}^i . Therefore, by definition of δ -better-response dynamics, we have that $\ell_k(\bar{t}_{-k}^i, t_k^{i+1}) \geq \ell_k(\bar{t}^i) + \delta$, and so by Lemmas A.7 and A.9, $\mu([t_k^{i+1}, t_k^i]) \geq \delta$ as required. \square

PROOF OF THEOREM 5.15. Let $\delta > 0$ and let $(\bar{t}^i, P_i)_{i=0}^\infty$ be a δ -better-response dynamic in (n, μ, \succeq_C) . By definition, $M(\bar{t}^0) \leq n$. By Lemma A.20, $M = M(\bar{t}^i)$ decreases by at least 1 in every $\lceil \frac{\mu(\mathcal{T})}{\delta n} \rceil$ rounds within which a Nash equilibrium is not reached. Therefore, if a Nash equilibrium is not reached in at most $n \cdot \lceil \frac{\mu(\mathcal{T})}{\delta n} \rceil$ rounds from 0, then $M(\bar{t}^i) = 0$ and so $\bar{t}^i \equiv 0$, which by Theorem 5.5 is a Nash equilibrium. If \bar{t}^i is a Nash equilibrium, then by Theorem 5.5, $M(\bar{t}^{i+1}) \leq n - 1$, and so if a Nash equilibrium is not reached in at most $(n - 1) \cdot \lceil \frac{\mu(\mathcal{T})}{\delta n} \rceil$ rounds from $i + 1$, then we have $M = 0$ once more, and so a Nash equilibrium is reached again.

The tighter bounds described in Remark 5.16 may be shown in a similar manner, due to Remark A.21. \square

Definition A.22 (*k*-Canonical Form). Let $k \in \mathbb{P}_n$ and let $\bar{t} \in \mathcal{T}^{\mathbb{P}_n}$ be a pure-strategy Nash equilibrium in (n, μ, \succeq_C) . We say that \bar{t} is in *k-canonical form* if all of the following hold.

- (1) $t_0 \leq t_1 \leq \dots \leq t_{k-1} \leq t_{k+1} \leq \dots \leq t_{n-1}$.
- (2) $\mu([0, t_j]) \leq \frac{j}{n}$ for every $j < k$.
- (3) Either $k = n - 1$, or $\mu([0, t_{k+1}]) > \frac{k}{n} \cdot \mu(\mathcal{T})$.

LEMMA A.23. Let $k \in \mathbb{P}_n$ and let $\bar{t} \in \mathcal{T}^{\mathbb{P}_n}$ be a pure-strategy Nash equilibrium in (n, μ, \succeq_C) . There exists a permutation $\pi \in \mathbb{P}_n!$ s.t. $(t_{\pi(0)}, t_{\pi(1)}, \dots, t_{\pi(n-1)})$ is in $\pi^{-1}(k)$ -canonical form.

PROOF. We start by defining π such that $t_{\pi(0)} \leq t_{\pi(1)} \leq \dots \leq t_{\pi(n-1)}$. In particular, we have

$$t_{\pi(0)} \leq t_{\pi(1)} \leq \dots \leq t_{\pi^{-1}(k)-1} \leq t_{\pi^{-1}(k)+1} \leq \dots \leq t_{\pi(n-1)}. \quad (1)$$

By Theorem 5.5, we have that

$$\mu([0, t_{\pi(j)}]) \leq \frac{j}{n} \cdot \mu(\mathcal{T}) \quad (2)$$

for every $j \in \mathbb{P}_n$, and in particular for every $j < \pi^{-1}(k)$. If $\pi^{-1}(k) = n - 1$ or $\mu([0, t_{\pi(\pi^{-1}(k)+1)}]) > \frac{\pi^{-1}(k)}{n} \cdot \mu(\mathcal{T})$, then π is a permutation as required. Otherwise, we have

$$\mu([0, t_{\pi(\pi^{-1}(k)+1)}]) \leq \frac{\pi^{-1}(k)}{n} \cdot \mu(\mathcal{T}). \quad (3)$$

In this case, we modify π to create a new permutation $\pi' \in \mathbb{P}_n!$ by incrementing $\pi^{-1}(k)$, or more formally — by swapping the values of coordinates $\pi^{-1}(k)$ and $\pi^{-1}(k) + 1$ of π . We note that Eq. (1) still holds w.r.t. π' (i.e. by substituting π' for π). By Eq. (2) w.r.t. π for all $j < \pi^{-1}(k)$, we have that Eq. (2) holds w.r.t. π' for all $j < \pi'^{-1}(k) - 1$; by Eq. (3) w.r.t. π , we have that Eq. (2) holds w.r.t. π' for $j = \pi'^{-1}(k) - 1$ as well. Once again, if $\pi'^{-1}(k) = n - 1$ or $\mu([0, t_{\pi'(\pi'^{-1}(k)+1)}]) > \frac{\pi'^{-1}(k)}{n} \cdot \mu(\mathcal{T})$, then π' is a permutation as required. Otherwise, Eq. (3) holds w.r.t. π' , and we repeat the modification step. As $\pi^{-1}(k)$ is incremented in each modification step, this process concludes in at most $n - 1$ steps, as it concludes if $\pi^{-1}(k)$ reaches $n - 1$. \square

LEMMA A.24. Let $k \in \mathbb{P}_n$ and let $\bar{t} \in \mathcal{T}^{\mathbb{P}_n}$ be a pure-strategy Nash equilibrium in (n, μ, \succeq_C) in *k-canonical form*. Both of the following hold.

- (1) $\mu([0, t_j]) \leq \frac{j}{n}$ for every $j \in \mathbb{P}_n$.
- (2) Either $k = n - 1$ or $t_k < t_{k+1}$.
- (3) For every $t \in \mathcal{T}$, (\bar{t}_{-k}, t) is a Nash equilibrium in (n, μ, \succeq_C) iff $\mu([0, t]) \leq \frac{k}{n}$.

PROOF. By definition of k -canonical form, $\mu([0, t_j]) \leq \frac{j}{n}$ for every $j < k$. By definition of k -canonical form, we also have for every $n > j > k$ that $\mu([0, t_j]) \geq \mu([0, t_{k+1}]) > \frac{k}{n} \cdot \mu(\mathcal{T})$. Therefore, by Theorem 5.5, $\mu([0, t_j]) \leq \frac{k}{n} \cdot \mu(\mathcal{T})$ for every $j \leq k$, and in particular $\mu([0, t_k]) \leq \frac{k}{n} \cdot \mu(\mathcal{T})$. Furthermore, we obtain that $(t_{k+1}, t_{k+2}, \dots, t_{n-1})$ are the $n - k - 1$ numerically-largest strategies in \bar{t} , and as they are sorted, by Theorem 5.5 we have that $\mu([0, t_j]) \leq \frac{j}{n} \cdot \mu(\mathcal{T})$ for every $j > k$ as well and the proof of Part 1 is complete.

Assume that $k < n - 1$. As we have shown that $\mu([0, t_k]) \leq \frac{k}{n} \cdot \mu(\mathcal{T})$, but by definition of k -canonical form $\mu([0, t_{k+1}]) > \frac{k}{n} \cdot \mu(\mathcal{T})$, we have that $t_k < t_{k+1}$ and Part 2 holds.

We conclude by proving Part 3; let $t \in \mathcal{T}$. If $\mu([0, t]) \leq \frac{k}{n}$, then by Part 1 and Theorem 5.5, (\bar{t}_{-k}, t) is a Nash equilibrium in (n, μ, \succeq_C) . Recall that $\mu([0, t_j]) > \frac{k}{n} \cdot \mu(\mathcal{T})$ for every $j > k$; therefore, if $\mu([0, t]) > \frac{k}{n} \cdot \mu(\mathcal{T})$ as well, then by Theorem 5.5, (\bar{t}_{-k}, t) is a not a Nash equilibrium in (n, μ, \succeq_C) . \square

PROOF OF THEOREM 5.18. Let $\bar{t} \in \mathcal{T}^{\mathbb{P}_n}$ be a Nash equilibrium in (n, μ, \succeq_C) and let $k \in \mathbb{P}_n$. By Lemma A.23, assume w.l.o.g. that \bar{t} is in k -canonical form. By Lemma A.24(3), it is enough to show that each $t_k < t'_k \leq 1$ s.t. $\mu([0, t'_k]) > \frac{k}{n} \cdot \mu(\mathcal{T})$ is not a best response to \bar{t}_{-k} in (n, μ, \succeq_C) . If $k < n + 1$, then by Lemma A.24(2), $t_k < t_{k+1}$ and so, by Lemma A.13 and since $\mu([0, t_{k+1}]) > \frac{k}{n} \cdot \mu(\mathcal{T})$, it is enough to consider $t_k < t'_k \leq t_{k+1}$ in this case. By definition, for every $j \leq k$, we have $\mu([0, t'_k]) > \frac{k}{n} \cdot \mu(\mathcal{T}) \geq \mu([0, t_j])$ and so $t'_k > t_j$.

Let s' and s'' be mixed-consumption Nash equilibria in the k -producer consumer game $(\mu|_{[0, t'_k]}; t_0, \dots, t_{k-1})$ and in the $(n-k)$ -producer game $(\mu|_{[t'_k, 1]}; t'_k, t_{k+1}, \dots, t_{n-1})$, respectively; by abuse of notation, we think of s'' as $s'' = (s''_{\neg}, s''_k, s''_{k+1}, \dots, s''_{n-1})$ ($s''_{\neg} \equiv 0$ by definition of s'') and for each $k \leq j < n$ define $\ell_j^{s''} \triangleq \int_{\mathcal{T}} s''_j d(\mu|_{[t'_k, 1]})$. For every $0 \leq j < k$, we have by definition of k -canonical form that

$$\mu([0, t_j]) \leq \frac{j}{n} \cdot \mu(\mathcal{T}) = \frac{j}{k} \cdot \frac{k}{n} \cdot \mu(\mathcal{T}) < \frac{j}{k} \cdot \mu([0, t'_k]) = \frac{j}{k} \cdot \mu|_{[0, t'_k]}(\mathcal{T}).$$

By Theorem 5.5, (t_0, \dots, t_{k-1}) is therefore a Nash equilibrium in $(k, \mu|_{[0, t'_k]}, \succeq_C)$, and so by Theorem 5.4 we have $\ell_j^{s'} = \frac{\mu|_{[0, t'_k]}(\mathcal{T})}{k} = \frac{\mu([0, t'_k])}{k} > \frac{\mu(\mathcal{T})}{n}$ for every $0 \leq j < k$.

For every $k < j < n$, we have by Lemma A.24(1) that

$$\mu([t_j, 1]) \geq \frac{n-j}{n} \cdot \mu(\mathcal{T}) > \frac{n-j}{n} \cdot \frac{n}{n-k} \cdot \mu([t'_k, 1]) = \frac{n-j}{n-k} \cdot \mu([t'_k, 1]),$$

and therefore

$$\mu([t'_k, t_j]) = \mu([t'_k, 1]) - \mu([t_j, 1]) < \mu([t'_k, 1]) - \frac{n-j}{n-k} \cdot \mu([t'_k, 1]) = \frac{j-k}{n-k} \cdot \mu|_{[t'_k, 1]}(\mathcal{T}).$$

Note that $\mu([t'_k, t'_k]) = 0 = \frac{k-k}{n-k} \cdot \mu|_{[t'_k, 1]}(\mathcal{T})$ trivially holds as well. By Theorem 5.5, $(t'_k, t_{k+1}, t_{k+2}, \dots, t_{n-1})$ is therefore a Nash equilibrium in $(n-k, \mu|_{[t'_k, 1]}, \succeq_C)$, and so by Theorem 5.4 we have that $\ell_j^{s''} = \frac{\mu|_{[t'_k, 1]}(\mathcal{T})}{n-k} = \frac{\mu([t'_k, 1])}{n-k} < \frac{\mu(\mathcal{T})}{n}$ for every $k \leq j < n$.

Let s be the mixed-consumption profile defined by $s_j|_{[0, t'_k]} = s'_j$ for every $j \in \{\neg, 0, 1, \dots, k-1\}$ and $s_j|_{[0, t'_k]} \equiv 0$ for every $k \leq j < n$, and by $s_j|_{[t'_k, 1]} = s''_j$ for every $k \leq j < n$ and $s_j|_{[t'_k, 1]} \equiv 0$ for every $j \in \{\neg, 0, 1, \dots, k-1\}$. By definition of s' and of s'' , we have that s is a legal mixed-consumption profile in $(\mu; \bar{t}_{-k}, t'_k)$, and furthermore, that $\ell_j^s = \ell_j^{s'} = \frac{\mu([0, t'_k])}{k}$ for every $j \in \{\neg, 0, 1, \dots, k-1\}$, and that $\ell_j^s = \ell_j^{s''} = \frac{\mu([t'_k, 1])}{n-k}$ for every $k \leq j < n$. By the former, and as s' is a Nash equilibrium, no type $d \in [0, t'_k)$

has any incentive to deviate from s in $(\mu; \bar{t}_{-k}, t'_k)$, and by the latter, as s'' is a Nash equilibrium and as $\frac{\mu([t'_k, 1])}{n-k} < \frac{\mu(\mathcal{T})}{n} < \frac{\mu([0, t'_k])}{k}$, we have that neither does any type $d \in [t'_k, 1]$. Therefore, s is a Nash equilibrium in $(\mu; \bar{t}_{-k}, t'_k)$. As $\ell_k(\bar{t}_{-k}, t'_k) = \ell_k^s = \frac{\mu([t'_k, 1])}{n-k} < \frac{\mu(\mathcal{T})}{n} = \ell_k(\bar{t})$ (with the last equality by Theorem 5.4, since \bar{t} is a Nash equilibrium in (n, μ, \succeq_C)), we have that producer k strictly prefers t_k over t'_k given \bar{t}_{-k} , and so t'_k is not a best response to \bar{t}_{-k} in (n, μ, \succeq_C) , as required.

We note that an alternative proof may also be given via Algorithm 1 / Corollary A.6. \square

PROOF OF COROLLARY 5.19. A direct corollary of Theorems 5.15 and 5.18 and Remark 5.17. \square

PROOF OF PROPOSITION 5.20. It is enough to show that some nonequilibrium can be reached in a finite number of steps from any Nash equilibrium. Let $t \in \mathcal{T}$ s.t. $0 < \mu([0, t]) \leq \frac{n-1}{n} \cdot \mu(\mathcal{T})$ (there must exist such t by definition of μ) and let $j \in \{1, \dots, n-1\}$ be minimal s.t. $\mu([0, t]) \leq \frac{j}{n} \cdot \mu(\mathcal{T})$.

By Theorem 5.3, $0 \in \mathcal{T}$ is a best-response by any producer to any Nash equilibrium, and so a (nonlazy) best-response dynamic can reach $(0, 0, \dots, 0)$ from any Nash equilibrium in one round. Let $i \in \mathbb{N}$ be more than one round into the future after reaching $(0, 0, \dots, 0)$, s.t. $|P_i| > 1$, and let $k, h \in P_i$ s.t. $k \neq h$. By definition of t and j and by Theorem 5.5, any strategy profile in which at most $n-j$ producers play t and the rest play 0 is a Nash equilibrium. Therefore, within one round after reaching $(0, 0, \dots, 0)$, a Nash equilibrium in which $j+1$ producers, including k and h , play 0 and the rest play t , can be reached, and can be lazily maintained until the step i . In step i , all triggered producers may play $t \in \mathcal{T}$, which is a best response for each of them since at least j producers playing 0 and the rest playing t is a Nash equilibrium. Therefore, and as k and h both switch from playing 0 to playing t at i , at least $n-j+1$ producers play t at $i+1$, which, by definition of j and by Theorem 5.5, is a nonequilibrium. \square

LEMMA A.25. *Let $k \in \mathbb{P}_n$, let $\bar{t}_{-k} \in \mathcal{T}^{\mathbb{P}_n \setminus \{k\}}$. If $\mu([0, t_j]) > 0$ for all $j \in \mathbb{P}_n \setminus \{k\}$, then $t'_k \in \mathcal{T}$ is a best response (by k) to \bar{t}_{-k} in (n, μ, \succeq_C) iff $\mu([0, t'_k]) = 0$.*

PROOF. By Lemma A.13, all strategies $t \in \mathcal{T}$ for which $\mu([0, t]) = 0$ are equivalent. As in particular, $t = 0 \in \mathcal{T}$ is such a strategy, it is therefore enough to show that k strictly prefers to play $0 \in \mathcal{T}$ over any t'_k s.t. $\mu([0, t'_k]) > 0$. By Lemma A.13(1), it is enough to consider the case in which $t'_k \leq t_j$ for all $j \in \mathbb{P}_n \setminus \{k\}$. By Algorithm 1 (for t_n as defined there), $\ell_k(\bar{t}_{-k}, 0) = \text{Max}_{0 < j \leq n} \frac{\mu([0, t_j])}{j} > \text{Max}_{0 < j \leq n} \frac{\mu([t'_k, t_j])}{j} = \ell_k(\bar{t}_{-k}, t'_k)$, as required. \square

PROOF OF THEOREM 5.22. As in the proof of Theorem 5.15, and as a best-response dynamic is also a $\mu(\mathcal{T})$ -better-response dynamic, we have that $M(\bar{t}^0) \leq n$ (and $M \leq n-1$ on every step immediately following a Nash equilibrium), and that M decreases by at least 1 every round if a Nash equilibrium is not reached. Let $i \in \mathbb{N}$ s.t. $M(\bar{t}^i) = 1$, $M(\bar{t}^{i-1}) = 2$, and \bar{t}^{i-1} is not a Nash equilibrium; it is enough to show that \bar{t}^i is a Nash equilibrium.

Since $M(\bar{t}^i) = 1$, by Theorem 5.5 it is enough to show that there exists $j \in \mathbb{P}_n$ s.t. $\mu([0, t_j^i]) = 0$. As $M(\bar{t}^i) \neq M(\bar{t}^{i-1})$, we have $P_{i-1} \neq \emptyset$. By Theorem 5.5, as $M(\bar{t}^{i-1}) = 2$ but \bar{t}^{i-1} is not a Nash equilibrium, there exists at most one producer $k \in \mathbb{P}_n$ s.t. $\mu([0, t_k^{i-1}]) = 0$. If there exists no such producer, then by Lemma A.25, we have $\mu([0, t_j^i]) = 0$ for every $j \in P_i$, and as $P_i \neq \emptyset$, the proof is complete. Otherwise, there exists a unique $k \in \mathbb{P}_n$ s.t. $\mu([0, t_k^{i-1}]) = 0$. If $k \notin P_i$, then $t_k^i = t_k^{i-1}$ and the proof is

complete. Otherwise, $k \in P_i$ and by Lemma A.25, we have $\mu([0, t_k^i]) = 0$, as required. \square

PROOF OF COROLLARY 5.25. A direct corollary of Theorems 5.18 and 5.22, Remark 5.17, and Example 5.24. \square

A.4. Proofs and Additional Results for Section 5.2

A.4.1. Proofs and Additional Results for Section 5.2.1

LEMMA A.26 (DOMINATION). *Let $t \neq t' \in \mathcal{T}$ be strategies in (n, μ, \succeq_F) . t is a safe alternative to t' iff either of the following hold. In either case, t strongly dominates t'*

- (1) $t > t'$ and $\mu([t', t]) = 0$.
- (2) $\mu([t', 1]) < \frac{\mu([t, 1])}{n}$.

PROOF. t is a safe alternative to (alternatively, strongly dominates) t' iff either t always produces greater load than t' , or t always produces at least as much load as t' and in addition $t > t'$. By Lemma A.13(3), the former occurs iff $\mu([t', 1]) < \frac{\mu([t, 1])}{n}$; by Lemma A.13(1,2), if $t > t'$, then the latter occurs iff $\mu([t', t]) = 0$ \square

LEMMA A.27. $\mu\left(\bigcup\left\{[t, t') \mid 0 \leq t < t' \ \& \ \mu([t, t']) \leq m\right\}\right) \leq m$, for every $t' \in \mathcal{T}$ and $m \in \mathbb{R}_{\geq}$.

PROOF. Define $U \triangleq \bigcup\left\{[t, t') \mid 0 \leq t < t' \ \& \ \mu([t, t']) \leq m\right\}$. If $U = \emptyset$, then $\mu(U) = 0 \leq m$ and the proof is complete; assume, therefore, that $U \neq \emptyset$ and let $u \triangleq \inf U \geq 0$. By definition, U is connected, and therefore by definition of U and u , we have that either $U = [u, t')$ or $U = (u, t')$. If $U = [u, t')$, then $u \in U$, and by definition of U , we obtain $\mu(U) = \mu([u, t')) \leq m$, as required; assume therefore, that $U = (u, t')$. In this case, $U = \bigcup\left\{[t, t') \mid t \in [0, t') \cap \mathbb{Q} \ \& \ \mu([t, t']) \leq m\right\}$, and by continuity of μ from below, we obtain $\mu(U) \leq m$, as required. \square

PROOF OF PROPOSITION 5.27. Let $t \in \mathcal{T}$ be a dominant strategy in this game. By Lemma A.26, both $\mu([0, t]) = 0$ (as t is a safe alternative to 0) and $\mu([t', 1]) < \frac{\mu([t, 1])}{n}$ for every $t' > t$ (as t is a safe alternative to every such t'). (Alternatively, the former holds as by definition, any strategy dominant w.r.t. fine preferences is also dominant w.r.t. coarse preferences, and by Theorem 5.3.) By the former, $\mu([t, 1]) = \mu(\mathcal{T})$, and therefore and by the latter, $\mu([t', 1]) < \frac{\mu(\mathcal{T})}{n}$ for every $t' \in (t, 1) \cap \mathbb{Q}$. Therefore, by continuity of μ from below, $\mu((t, 1]) \leq \frac{\mu(\mathcal{T})}{n}$. Hence, $\mu(\{t\}) = \mu([0, t]) \geq \frac{n-1}{n} \cdot \mu(\mathcal{T})$. (Conversely, by Lemma A.26, it is easy to see that if there indeed exists $t \in \mathcal{T}$ s.t. $\mu([0, t]) = 0$ and $\mu([t', 1]) < \frac{\mu(\mathcal{T})}{n}$ for every $t' > t$, then it constitutes the unique dominant strategy.)

For the second statement, we note that by Lemma A.26, the set of dominated strategies is

$$\begin{aligned} & \left\{t \in \mathcal{T} \mid \mu([t, 1]) < \frac{\mu(\mathcal{T})}{n}\right\} \cup \left\{t \in \mathcal{T} \mid \exists t < t' \leq 1 : \mu([t, t']) = 0\right\} = \\ & = \bigcup\left\{[t, 1] \mid t \in \mathcal{T} \ \& \ \mu([t, 1]) < \frac{\mu(\mathcal{T})}{n}\right\} \cup \bigcup\left\{[t, t') \mid t' \in \mathcal{T} \cap \mathbb{Q} \ \& \ 0 \leq t < t' \ \& \ \mu([t, t']) = 0\right\}. \end{aligned}$$

By σ -additivity of μ and by Lemma A.27 (applied twice),

$$\begin{aligned}
& \mu\left(\left\{t \in \mathcal{T} \mid \mu([t, 1]) < \frac{\mu(\mathcal{T})}{n}\right\} \cup \left\{t \in \mathcal{T} \mid \exists t' < t' \leq 1 : \mu([t, t']) = 0\right\}\right) \leq \\
& \leq \mu\left(\bigcup\left\{[t, 1] \mid t \in \mathcal{T} \ \& \ \mu([t, 1]) < \frac{\mu(\mathcal{T})}{n}\right\}\right) + \\
& \qquad \qquad \qquad + \sum_{t' \in \mathcal{T} \cap \mathbb{Q}} \mu\left(\bigcup\left\{[t, t'] \mid 0 \leq t < t' \ \& \ \mu([t, t']) = 0\right\}\right) \leq \\
& \leq \frac{\mu(\mathcal{T})}{n} + 0 = \frac{\mu(\mathcal{T})}{n}.
\end{aligned}$$

We note that this bound is attained if μ is atomless, as in this case it is straightforward to verify that $\mu(\bigcup\{[t, 1] \mid t \in \mathcal{T} \ \& \ \mu([t, 1]) < \frac{\mu(\mathcal{T})}{n}\}) = \frac{\mu(\mathcal{T})}{n}$.

The second statement leads to an extremely concise, yet somewhat more obscure, proof for the first one. If a dominant strategy exists, then it is a safe alternative to all other strategies; in particular, all other strategies have safe alternatives (other than themselves). By the second statement, at least $\frac{n-1}{n}$ of μ is therefore concentrated on this dominated strategy, and the proof is complete. \square

LEMMA A.28. $\{t \in \mathcal{T} \mid \mu([t', t]) \leq m\}$ attains a maximum value for every $t' \in \mathcal{T}$ and $m \in \mathbb{R}_{\geq}$.

PROOF. Denote $S \triangleq \{t \in \mathcal{T} \mid \mu([t', t]) \leq m\}$. We note that $t' \in S$. If $t' = \sup S$, then $t' = \text{Max } S$ and the proof is complete. Assume, therefore, that $t' < \sup S$. Let $U \triangleq \bigcup\{[t', t] \mid t \in \mathcal{T} \ \& \ \mu([t', t]) \leq m\}$. As $t' < \sup S$, we have that $U \neq \emptyset$. Let $u \triangleq \sup U \leq 1$. By definition of U and of u , we have $U = [t', u)$. By definition of U and of S , we have that $u = \sup U = \sup S$, and so it is enough to show that $u \in S$, i.e. that $\mu([t', u]) \leq m$. As $U = [t', u)$, this is equivalent to showing that $\mu(U) \leq m$. Observe that

$$U = \bigcup\left\{[t', t] \mid t \in \mathcal{T} \cap \mathbb{Q} \ \& \ \mu([t', t]) \leq m\right\}.$$

By continuity of μ from below, we thus obtain $\mu(U) \leq m$, as required. \square

PROOF OF THEOREM 5.28. For every $j \in \mathbb{P}_n$, let $t_j \triangleq \text{Max}\{t \in \mathcal{T} \mid \mu([0, t]) \leq \frac{j}{n} \cdot \mu(\mathcal{T})\}$. (t_j is well-defined by Lemma A.28.) By Theorem 5.5, \bar{t} is a Nash equilibrium in (n, μ, \succeq_C) , and so by Theorem 5.4, $\ell_j(\bar{t}) = \frac{\mu(\mathcal{T})}{n}$ for every $j \in \mathbb{P}_n$. We now first show that no Nash equilibrium other than \bar{t} (up to permutations) exists in (n, μ, \succeq_F) , and then show that \bar{t} (and hence all permutations thereof) is a super strong Nash equilibrium in (n, μ, \succeq_F) .

Let $t'_0 \leq \dots \leq t'_{n-1} \in \mathcal{T}$ s.t. \bar{t}' is a Nash equilibrium in (n, μ, \succeq_F) . We will show that $t'_j = t_j$ for every $j \in \mathbb{P}_n$. By definition of coarse and fine preferences, \bar{t}' is also a Nash equilibrium in (n, μ, \succeq_C) . Therefore, by Theorem 5.5 we have that $\mu([0, t'_j]) \leq \frac{j}{n} \cdot \mu(\mathcal{T})$ for every $j \in \mathbb{P}_n$. Hence we have for every $j \in \mathbb{P}_n$ both that $t'_j \leq t_j$ and (by Theorem 5.5 again) that (\bar{t}'_{-j}, t'_j) is a Nash equilibrium in (n, μ, \succeq_C) as well. Therefore, by Theorem 5.4, $\ell_j(\bar{t}') = \frac{\mu(\mathcal{T})}{n} = \ell_j(\bar{t}'_{-j}, t'_j)$. As \bar{t}' is a Nash equilibrium in (n, μ, \succeq_F) , we therefore have that $t'_j \geq t_j$, and so $t'_j = t_j$, as required.

We now show that \bar{t} is a super strong Nash equilibrium in (n, μ, \succeq_F) . Assume for contradiction that there exists a coalition $P \subseteq \mathbb{P}_n$ and strategies $\bar{t}' = (t'_j)_{j \in P} \in \mathcal{T}^P$ s.t. j weakly prefers (\bar{t}_{-P}, \bar{t}') over \bar{t} w.r.t. fine preferences for every $j \in P$, with a strict preference for at least one producer $j \in P$. For every $j \in P$, as j weakly prefers (\bar{t}_{-P}, \bar{t}') over \bar{t} , we have that $\ell_j(\bar{t}_{-P}, \bar{t}') \geq \ell_j(\bar{t})$. As \bar{t} is a Nash equilibrium in (n, μ, \succeq_C) , by Theorem 5.6, we therefore have $\ell_j(\bar{t}_{-P}, \bar{t}') = \ell_j(\bar{t})$ for every $j \in P$. Therefore, by definition

of P and \bar{t}' , we have $t'_j \geq t_j$ for every $j \in P$, with a strict inequality for at least one producer $j \in P$ — let \hat{j} be such a producer for which $t'_{\hat{j}}$ is greatest. Assume w.l.o.g. that either $\hat{j} = n - 1$ or $t_{\hat{j}} < t_{\hat{j}+1}$; therefore, \bar{t} is in \hat{j} -canonical form. As \bar{t} is also a Nash equilibrium in (n, μ, \succeq_C) , by Theorem 5.18, by Lemma A.24(3), by definition of $t_{\hat{j}}$, and as $t'_{\hat{j}} > t_{\hat{j}}$, we have $\ell_{\hat{j}}(\bar{t}) > \ell_{\hat{j}}(\bar{t}_{-\hat{j}}, t'_{\hat{j}})$. By Lemma A.8 and by definition of \hat{j} , we have $\ell_{\hat{j}}(\bar{t}_{-\hat{j}}, t'_{\hat{j}}) \geq \ell_{\hat{j}}(\bar{t}_{-P}, \bar{t}')$. Therefore, $\ell_{\hat{j}}(\bar{t}) > \ell_{\hat{j}}(\bar{t}_{-\hat{j}}, t'_{\hat{j}}) \geq \ell_{\hat{j}}(\bar{t}_{-P}, \bar{t}')$ — a contradiction. \square

PROOF OF COROLLARY 5.29. Since the CDF of μ is continuous and strictly increasing, for every $j \in \mathbb{P}_n$ there exists a unique strategy $t_j \in \mathcal{T}$ s.t. $\mu([0, t_j]) = \frac{j}{n} \cdot \mu(\mathcal{T})$; hence, $t_j = \text{Max}\{t \in \mathcal{T} \mid \mu([0, t]) \leq \frac{j}{n} \cdot \mu(\mathcal{T})\}$ and by Theorem 5.28 the proof is complete. \square

PROOF OF PROPOSITION 5.30. Direct from definition of (n, μ, \succeq_F) , as no player is ever indifferent between any two strategies, regardless of the information such player possesses regarding the strategies of the other players. \square

PROOF OF THEOREM 5.31. By Theorem 5.28, for every $j \in \mathbb{P}_n$ we have

$$\{d \in \mathcal{T} \mid \mu([0, d]) \in (\frac{j}{n} \cdot \mu(\mathcal{T}), \frac{j+1}{n} \cdot \mu(\mathcal{T}))\} \subseteq [t_j, t_{j+1}).$$

Therefore, it is enough to show that $s_j(d) = 1$ for almost all $d \in [t_j, t_{j+1})$. By definition of s , this is equivalent to showing that $\int_{[t_j, t_{j+1})} s_j d\mu = \mu([t_j, t_{j+1}))$ for every $j \in \mathbb{P}_n$, where $t_n \triangleq 1$. As μ is atomless, by Theorem 5.28 and by definition of t_n we have $\mu([t_j, t_{j+1})) = \mu([0, t_{j+1})) - \mu([0, t_j]) = \frac{j+1}{n} \cdot \mu(\mathcal{T}) - \frac{j}{n} \cdot \mu(\mathcal{T}) = \frac{\mu(\mathcal{T})}{n}$ for every $j \in \mathbb{P}_n$. We prove the theorem by induction on j . Let $k \in \mathbb{P}_n$ and assume that the theorem holds for every $0 \leq j < k$; we now show that it holds for $j = k$ as well.

For every $0 \leq j < k$, by the induction hypothesis, $\int_{[t_j, t_{j+1})} s_j d\mu = \mu([t_j, t_{j+1})) = \frac{\mu(\mathcal{T})}{n}$. By Theorem 5.28 and by definition of $\ell_j(\bar{t})$ and of ℓ_j^s ,

$$\frac{\mu(\mathcal{T})}{n} = \ell_j(\bar{t}) = \ell_j^s = \int_{\mathcal{T}} s_j d\mu \geq \int_{[t_j, t_{j+1})} s_j d\mu + \int_{[t_k, t_{k+1})} s_j d\mu = \frac{\mu(\mathcal{T})}{n} + \int_{[t_k, t_{k+1})} s_j d\mu,$$

and so $\int_{[t_k, t_{k+1})} s_j d\mu = 0$. By definition of Nash equilibrium, $s_{-j}(d) = 0$ for every $d \geq t_k$, and therefore $\int_{[t_k, t_{k+1})} s_{-j} d\mu = 0$ as well. Let $S_k \triangleq \{-, 0, 1, 2, \dots, k-1\}$; by definition of s , we have that $s(d) \in \Delta^{S_k}$ for every $d \in [t_k, t_{k+1})$. Therefore,

$$\int_{[t_k, t_{k+1})} s_k d\mu = \mu([t_k, t_{k+1})) - \sum_{j \in S_k} \int_{[t_k, t_{k+1})} s_j d\mu = \frac{\mu(\mathcal{T})}{n} - 0 = \mu([t_k, t_{k+1})),$$

and the proof by induction is complete. \square

A.4.2. Proofs for Section 5.2.2

PROOF OF PROPOSITION 5.34 (NONCONSTRUCTIVE). Let $\tilde{\ell} \triangleq \ell_j(\bar{t}_{-j}, 0)$; by Theorem 5.3, $\tilde{\ell}$ is the maximum load attainable by j given \bar{t}_{-j} . Define $S \triangleq \{\mu([0, t]) \mid t \in \mathcal{T} \ \& \ \ell_j(\bar{t}_{-j}, t) = \tilde{\ell}\}$. Observe that $S \neq \emptyset$ as $0 \in S$ (given by $t = 0$); let $m \triangleq \sup S$. By Lemma A.13(1), every $t \in \mathcal{T}$ s.t. $\mu([0, t]) < m$ maximizes $\ell_j(\bar{t}_{-j}, t)$, while every $t \in \mathcal{T}$ s.t. $\mu([0, t]) > m$ does not. Assume for contradiction that there exists $t \in \mathcal{T}$ s.t. $\mu([0, t]) = m$ and $\ell_j(\bar{t}_{-j}, t) < \tilde{\ell}$. Let $\varepsilon \triangleq \tilde{\ell} - \ell_j(\bar{t}_{-j}, t) > 0$. By definition of m , there exists t' s.t. $m \geq \mu([0, t']) > m - \varepsilon$ and $\ell_j(\bar{t}_{-j}, t') = \tilde{\ell} = \ell_j(\bar{t}_{-j}, t) + \varepsilon$; by Lemma A.9, this is a contradiction. (We note that we have not shown (yet) that there exists $t \in \mathcal{T}$ s.t. $\mu([0, t]) = m$, but rather that every such t maximizes the load on j .) Therefore, we have

that the set of load-maximizing strategies for j is precisely $\{t \in \mathcal{T} \mid \mu([0, t]) \leq m\}$. By Lemma A.28, this set attains a maximum value. As by definition we have that a best response in (n, μ, \succeq_F) is a numerically-largest load-maximizing response, we obtain that this maximum value is a best response as required. Uniqueness follows directly from definition of fine preferences. \square

Before constructively proving Proposition 5.34, we first constructively prove it for two special cases.

COROLLARY A.29 (PROPOSITION 5.34 — SPECIAL CASE: LARGE STRATEGIES). *Let $k \in \mathbb{P}_n$, and let $\bar{t}_{-k} \in \mathcal{T}^{\mathbb{P}_n \setminus \{k\}}$. If $\mu([0, t_j]) > 0$ for all $j \in \mathbb{P}_n \setminus \{k\}$, then the unique best response (by k) to \bar{t}_{-k} in (n, μ, \succeq_F) is $\text{Max}\{t \in \mathcal{T} \mid \mu([0, t]) = 0\}$.*

PROOF. A direct corollary of Lemma A.25, as a best response in (n, μ, \succeq_F) is a numerically-largest load-maximizing response; the specified strategy is well defined by Lemma A.28. \square

LEMMA A.30 (PROPOSITION 5.34 — SPECIAL CASE: COARSE EQUILIBRIUM). *Let $\bar{t} \in \mathcal{T}^{\mathbb{P}_n}$ be a Nash equilibrium in (n, μ, \succeq_C) . For every $j \in \mathbb{P}_n$, a best response (by j) to \bar{t}_{-j} exists in (n, μ, \succeq_F) .*

PROOF (CONSTRUCTIVE). By Lemma A.23, assume w.l.o.g. that \bar{t} is in j -canonical form. We will show that a best response as required is given by $t'_j \triangleq \text{Max}\{t \in \mathcal{T} \mid \mu([0, t]) \leq \frac{j}{n} \cdot \mu(\mathcal{T})\}$. (t'_j is well defined by Lemma A.28.) By Lemma A.24(3) and by Theorem 5.18, a strategy $t \in \mathcal{T}$ maximizes $\ell_j(\bar{t}_{-j}, t)$ iff $\mu([0, t]) \leq \frac{j}{n} \cdot \mu(\mathcal{T})$. As by definition we have that a best response in (n, μ, \succeq_F) is a numerically-largest load-maximizing response, we obtain that t'_j is a best response as required. \square

PROOF OF PROPOSITION 5.34 (CONSTRUCTIVE). W.l.o.g. we prove the result for $j = 0$. Assume w.l.o.g. that $t_1 \leq t_2 \leq \dots \leq t_{n-1}$. Uniqueness follows directly from definition of fine preferences; it is therefore enough to show that a best response exists. If $\mu([0, t_1]) > 0$, then by Corollary A.29 a best response exists as required. Assume therefore henceforth that $\mu([0, t_1]) = 0$. If $(\bar{t}_{-0}, 0)$ is a Nash equilibrium in (n, μ, \succeq_C) , then by Lemma A.30 a best response exists as required. Assume therefore henceforth that $(\bar{t}_{-0}, 0)$ is not a Nash equilibrium in (n, μ, \succeq_C) .

Let $k \in \mathbb{P}_n$ be minimal s.t. $\ell_k(\bar{t}_{-0}, 0) < \ell_0(\bar{t}_{-0}, 0)$. (Such k exists by Corollary A.15, since $(\bar{t}_{-0}, 0)$ is not a Nash equilibrium in (n, μ, \succeq_C) ; by definition, $k > 0$.) By definition of k and by Corollary A.4 and Lemma A.14, $(0, t_1, t_2, \dots, t_{k-1})$ is a Nash equilibrium in $(k, \mu|_{\cap[0, t_k]}, \succeq_C)$. By Lemma A.30, there exists a best response $t_0 \in \mathcal{T}$ to $(t_1, t_2, \dots, t_{k-1})$ in $(k, \mu|_{\cap[0, t_k]}, \succeq_F)$. We claim that t_0 is a best response to \bar{t}_{-0} in (n, μ, \succeq_F) as well.

As $\ell_k(\bar{t}_{-0}, 0) < \ell_0(\bar{t}_{-0}, 0)$, we have that $\ell_0(\bar{t}_{-0}, 0) \leq \mu([0, t_k])$, and in particular $\mu([0, t_k]) > 0$. Therefore, by Lemma A.13(3) and by definition of t_0 , we obtain $t_0 < t_k$. By Theorem 5.18, $(t_0, t_1, \dots, t_{k-1})$ is also a Nash equilibrium in $(k, \mu|_{\cap[0, t_k]}, \succeq_C)$. Therefore, by Theorem 5.4, $\ell_j(\mu|_{\cap[0, t_k]}; 0, t_1, t_2, \dots, t_{k-1}) = \frac{\mu([0, t_k])}{k} = \ell_j(\mu|_{\cap[0, t_k]}; t_0, t_1, \dots, t_{k-1})$ for every $j \in \mathbb{P}_k$. By the construction in the proof of Theorem 4.10 and as $t_0 < t_k$, therefore $\ell_j(\bar{t}) = \ell_j(\bar{t}_{-0}, 0)$ for every $j \in \mathbb{P}_n$. By Corollary A.4 and by definition of k , we have $\ell_j(\bar{t}) = \ell_j(\bar{t}_{-0}, 0) = \ell_j(\mu|_{\cap[0, t_k]}; 0, t_1, t_2, \dots, t_{k-1}) = \frac{\mu([0, t_k])}{k}$ for every $j \in \mathbb{P}_k$. As $\ell_0(\bar{t}) = \ell_0(\bar{t}_{-0}, 0)$, by Theorem 5.3 we have that t_0 maximizes the load on producer 0 in (n, μ, \succeq_F) .

Let $h \in \mathbb{P}_k$ s.t. $t_h \leq t_0 < t_{h+1}$. Such $h > 0$ exists as $\mu([0, t_1]) = 0$ and since $t_0 \geq \text{Max}\{t \in \mathcal{T} \mid \mu([0, t]) = 0\} \geq t_1$ by Lemma A.13 (this maximum value is attained by

Lemma A.28), and $h < k$ since $t_0 < t_k$. It remains to show that every $t'_0 \in (t_0, 1]$ does not maximize the load on producer 0 in (n, μ, \succeq_F) . By Lemma A.13(1), it is enough to consider the case $t_0 < t'_0 < t_{h+1}$. Note that $t_k > t_0 \geq t_1$ and so $k > 1$.

By definition of t'_0 and t_0 , we have that t'_0 does not maximize the load on producer 0 in $(k, \mu|_{[0, t_k]}, \succeq_F)$. By Corollary A.4 and by Corollary A.15 (since $\mu([0, t_1]) = 0$), $\ell_1(\bar{t}_{-0}, t'_0) \geq \ell_1(\mu|_{[0, t_k]}; t'_0, t_1, \dots, t_{k-1}) > \frac{\mu([0, t_k])}{k} = \ell_1(\bar{t})$, and so by Algorithm 1, $\ell_1(\bar{t}_{-0}, t'_0) = \frac{\mu([0, t'_0])}{h}$ and so $\ell_j(\bar{t}_{-0}, t'_0) = \frac{\mu([0, t'_0])}{h}$ for every $0 < j \leq h$, and in particular for $j = h$. Therefore, and by Lemma A.13(1), $\ell_h(\bar{t}_{-0}, t'_0) = \ell_1(\bar{t}_{-0}, t'_0) > \ell_1(\bar{t}) = \ell_0(\bar{t}) \geq \ell_0(\bar{t}_{-0}, t'_0)$. As $\frac{\mu([0, t'_0])}{h} = \ell_1(\bar{t}_{-0}, t'_0) > \frac{\mu([0, t_k])}{k}$, we have $\mu([t'_0, t_j]) = \mu([0, t_j]) - \mu([0, t'_0]) < \mu([0, t_j]) - \frac{h}{k} \cdot \mu([0, t_k])$ for every $j > h$. As $\ell_h(\bar{t}_{-0}, t'_0) > \ell_0(\bar{t}_{-0}, t'_0)$, by definition of h and by Algorithm 1 and Corollary A.6, we obtain (for t_n as defined there) that $\ell_0(\bar{t}_{-0}, t'_0) = \text{Max}_{h < j \leq n} \frac{\mu([t'_0, t_j])}{j-h} < \text{Max}_{h < j \leq n} \left(\mu([0, t_j]) - \frac{h}{k} \cdot \mu([0, t_k]) \right) \cdot \frac{1}{j-h} = \text{Max}_{h < j \leq n} \left(\mu([0, t_j]) - \sum_{i=1}^h \ell_i(\bar{t}) \right) \cdot \frac{1}{j-h} = \ell_0(\bar{t})$, and the proof is complete. \square

PROOF OF COROLLARY 5.35. A direct corollary of Theorems 5.15 and 5.22, as any weakly-/ δ -better-/best-response dynamic w.r.t. fine preferences is also a weakly-/ δ -better-/best-response dynamic w.r.t. coarse preferences. \square

PROOF OF THEOREM 5.36. Let $(\bar{t}^0, P_i)_{i=0}^\infty$ be a sequential δ -better-response dynamic in (n, μ, \succeq_F) s.t. \bar{t}^0 is a Nash equilibrium w.r.t. (n, μ, \succeq_C) . Let $k \in \mathbb{P}_n$ and let \tilde{i} be minimal s.t. $k \in P_{\tilde{i}}$. It is enough to show that t_k^i is constant for $i > \tilde{i}$, as this implies that after one round \bar{t}^i is constant regardless of P_i , and is thus a Nash equilibrium in (n, μ, \succeq_F) .

By Theorem 5.18, \bar{t}^i is a Nash equilibrium w.r.t. (n, μ, \succeq_C) for every $i \in \mathbb{N}$; therefore, by Theorem 5.4, the loads on all producers are constant throughout this dynamic. Therefore, by definition of δ -better-response dynamics, we have both that $(\bar{t}^0, P_i)_{i=0}^\infty$ is a best-response dynamic in (n, μ, \succeq_F) , and that $(t_j^i)_{i=0}^\infty$ is monotone-nondecreasing for every $j \in \mathbb{P}_n$.

Let $i > \tilde{i}$ s.t. $k \in P_i$. As $t_k^{i+1} \geq t_k^{\tilde{i}+1}$, it is enough to show that $t_k^{i+1} \leq t_k^{\tilde{i}+1}$. As $\bar{t}^{i+1} \geq (\bar{t}_{-k}^{\tilde{i}+1}, t_k^{i+1})$ in every coordinate, by Theorem 5.5 and since $\bar{t}^{\tilde{i}+1}$ is a Nash equilibrium w.r.t. (n, μ, \succeq_C) , so is $(\bar{t}_{-k}^{\tilde{i}+1}, t_k^{i+1})$, and so t_k^{i+1} maximizes the load on k given $\bar{t}_{-k}^{\tilde{i}+1}$; therefore, and as $t_k^{\tilde{i}+1}$ is a best response to $\bar{t}_{-k}^{\tilde{i}+1}$ w.r.t. (n, μ, \succeq_F) , we have that $t_k^{i+1} \leq t_k^{\tilde{i}+1}$, and the proof is complete. \square

PROOF OF COROLLARY 5.37. A direct corollary of Corollary 5.35 (Theorem 5.15) and Theorem 5.36. \square

PROOF OF COROLLARY 5.38. A direct corollary of Corollary 5.35 (Theorem 5.22) and Theorem 5.36. \square

A.5. Proof of Theorem 7.1

PROOF OF THEOREM 7.1. The fact that each such strategy profile is a super-strong equilibrium with load ℓ_j^g on each producer $j \in \mathbb{P}_{n_g}$ of good $g \in \{1, 2\}$ (and with the market allocated between the producers of each good as in the one-good scenario of Section 6) is an immediate consequence of Theorem 6.2; we therefore show that no other super-strong equilibrium exists. We give a proof for atomless μ ; the proof for general μ is similar and is left to the reader. Let $((t_j^g, \theta_j^g))_{j \in \mathbb{P}_{n_g}}^{g \in \{1, 2\}}$ be a super-strong equilibrium among producers, given in polar coordinates. For every $g \in \{1, 2\}$ and

$j \in \mathbb{P}_{n_g}$, let ℓ_j^g be the load on producer j of good g in $((t_j^g, \theta_j^g))_{j \in \mathbb{P}_{n_g}}^{g \in \{1,2\}}$.

We begin by noting that a producer $j \in \mathbb{P}_{n_g}$ of good $g \in \{1, 2\}$ may still secure a load of at least $\tilde{\ell}_j^g$ by choosing the origin as its location, and so $\ell_j^g \geq \tilde{\ell}_j^g$ for every $g \in \{1, 2\}$ and $j \in \mathbb{P}_{n_g}$. As for every $g \in \{1, 2\}$, we have $\sum_{j \in \mathbb{P}_{n_g}} \ell_j^g \leq \mu(\mathcal{T}) = \sum_{j \in \mathbb{P}_{n_g}} \tilde{\ell}_j^g$, we have that $\ell_j^g = \tilde{\ell}_j^g$ for every $j \in \mathbb{P}_{n_g}$, i.e. the load on every producer is as in the super-strong equilibria described in the statement of the theorem.

For every $g \in \{1, 2\}$, let $\pi_g \in \mathbb{P}_{n_g}!$ be a permutation s.t. $t_{\pi_g(0)}^g \leq t_{\pi_g(1)}^g \leq \dots \leq t_{\pi_g(n-1)}^g$. For every $g \in \{1, 2\}$ and $j \in \mathbb{P}_{n_g}$, define $m_j^g \triangleq \sum_{k=0}^{j-1} \tilde{\ell}_{\pi_g(k)}^g$; by the genericity assumption on partial sums of producer-equilibrium loads, and by positivity of equilibrium loads, we have $m_j^1 \neq m_k^2$ for every $j \in \mathbb{P}_{n_1}$ and $k \in \mathbb{P}_{n_2}$ s.t. either $j > 0$ or $k > 0$.

As for every $j \in \mathbb{P}_{n_g}$, the distance $t_{\pi_g(j)}^g$ is accessible by at least all consumer types consuming a positive amount from any of the producers $\pi_g(j), \pi_g(j+1), \dots, \pi_g(n-1)$ of good g , we have that $\mu([0, t_{\pi_g(j)}^g]) \leq \sum_{k=0}^{j-1} \tilde{\ell}_{\pi_g(k)}^g = m_j^g$ for every $j \in \mathbb{P}_{n_g}$. Therefore, deviating to a super-strong equilibrium as in the statement of the theorem, while maintaining the order of distances from the origin among producers of the same good, harms no producer. If $t_{\pi_g(j)}^g < \text{Max}\{t \in \mathcal{T} \mid \mu([0, t]) \leq m_j^g\}$ for some $g \in \{1, 2\}$ and $j \in \mathbb{P}_{n_g}$, then producer $\pi_g(j)$ of load g strictly benefits from such a deviation; therefore, $t_{\pi_g(j)}^g = \text{Max}\{t \in \mathcal{T} \mid \mu([0, t]) \leq m_j^g\}$ for every $g \in \{1, 2\}$ and $j \in \mathbb{P}_{n_g}$. Therefore, as μ is atomless, the market allocated between the producers of each good as in the one-good scenario of Section 6.

Assume for contradiction that not all producer strategies in $((t_j^g, \theta_j^g))_{j \in \mathbb{P}_{n_g}}^{g \in \{1,2\}}$ lie on the same ray from the origin. Therefore, w.l.o.g. there exist producers $j \in \mathbb{P}_{n_1}$ and $k \in \mathbb{P}_{n_2}$ whose strategies do not lie on the same ray, s.t. $t_j^1 \leq t_k^2$ and either $j = \pi_1(n_1 - 1)$ or $t_k^2 \leq t_{\pi_1(\pi_1^{-1}(j)+1)}^1$. (The w.l.o.g. assumption refers to the part played by each good.) If $j < \pi_1(n_1 - 1)$, then as μ is atomless, we have $\mu([0, t_k^2]) = m_{\pi_2^{-1}(k)}^2 \neq m_{\pi_1^{-1}(j)+1}^1 = \mu([0, t_{\pi_1(\pi_1^{-1}(j)+1)}^1])$ and so $t_k^2 < t_{\pi_1(\pi_1^{-1}(j)+1)}^1$; otherwise, since by assumption $\tilde{\ell}_k^2 > 0$ and as μ is atomless, we have $t_k^2 < 1$. Either way, by market allocation there exists $\varepsilon > 0$ s.t. almost all (w.r.t. μ) consumer types $d \in [t_k^2, t_k^2 + \varepsilon)$ consume a positive amount both from producer j of good 1 and from producer k of good 2. Let c be the circumference of the triangle whose vertices are the origin, (t_j^1, θ_j^1) and (t_k^2, θ_k^2) ; as the latter two do not lie on the same ray from the origin, (t_j^1, θ_j^1) is not a convex combination of the origin and (t_k^2, θ_k^2) , and so by the triangle inequality we have $c > 2t_k^2$. By definition of c , no consumer with type $d \in [t_k^2, \frac{c}{2})$ can consume from both producer j of good 1 and producer k of good 2 without violating the consumer's QoS limit. By combining these two, we have that for $\delta \triangleq \min\{\frac{c}{2} - t_k^2, \varepsilon\} > 0$, almost all consumer types $d \in [t_k^2, t_k^2 + \delta)$ consume from both these producers, while no such consumer consumes from both of them — a contradiction, since by definition of t_k^2 we have that $\delta > 0$ implies $\mu([t_k^2, t_k^2 + \delta)) > 0$. \square

We note that the requirements in Theorem 7.1, both for every producer-equilibrium load to be positive and for the genericity of partial sums of producer-equilibrium loads, are required. Indeed, any producer with zero producer-equilibrium load can be moved to any ray without destabilizing the equilibrium. Furthermore, if there exist permutations $\pi_1 \in \mathbb{P}_{n_1}!$ and $\pi_2 \in \mathbb{P}_{n_2}!$ and producers $j \in \mathbb{P}_{n_1} \setminus \{0\}$ and $k \in \mathbb{P}_{n_2} \setminus \{0\}$ s.t. $\sum_{i=0}^{j-1} \tilde{\ell}_{\pi_1(i)}^1 = \sum_{i=0}^{k-1} \tilde{\ell}_{\pi_2(i)}^2$, then moving all producers j' of good 1 s.t. $\pi_1(j') \geq \pi_1(j)$ and all producers k' of good 2 s.t. $\pi_2(k') \geq \pi_2(k)$ together to any ray does not destabilize the equilibrium either.