

Chapter 5

Competitive Market Model

The competitive market model serves as the basis for the two different multi-user allocation methods presented in this thesis. This market model prices resources based on supply and demand. Consumers (users) purchase these resources at the market price in order to maximize their happiness (QoS for a network application). Producers (switches) own the resources and maximize their utility by selling or renting. This model was chosen because of its ability to achieve certain desirable goals, such as Pareto-optimal distribution and price stability. The competitive market also has a structure that is simple to implement, and a well founded mathematical basis for analysis.

In this chapter, a description of the competitive market is provided as well as definitions of optimality (Pareto) and fairness (weighted max-min and equitable). Proofs that an economic model, consisting of multiple competitive markets, can achieve these measures of fairness are then given. An algorithm that determines the wealth distribution required for an equitable allocation is provided. Finally, some examples of optimal allocations are given for a simple network economy.

5.1 Market Definition

The competitive market model consists of scarce resources and two types of agents, consumers and producers. A resource is an item (or service) which is valued by agents in the economy. Since it is scarce, there is never enough of the resource to satisfy all the agents all

the time. For this reason, allocation decisions must be made. The agents come together at a market, where they buy or sell resources. Usually these exchanges are intermediated with money and the exchange rate of a resource is called its price. In the competitive market, prices are adjusted until supply equals demand. At this price the market is in *equilibrium* and the resulting allocation is Pareto-optimal [100]. The economies presented in this thesis will consist of multiple independent competitive markets, where each resource type will be sold in its own market.

Consumer j has wealth w^j and acts independently (selfishly) purchasing resources to increase utility. For each resource, it is assumed that the utility function, of each user, is monotonically increasing [100]. In addition, a user normally becomes satiated with some amount, above which the utility may decrease¹. Assume consumer j desires a maximum resource amount b^j ; therefore $u^j(b^j)$ is the highest utility consumer j can achieve. When maximizing utility, consumers must adhere to their budget constraints. Assuming consumer j wishes to purchase an amount a^j , where $a^j \leq b^j$, at price p the budget constraint $p \cdot a^j \leq w^j$ must be true. The wealth signifies purchasing power of each consumer, since consumers with more wealth can afford more resources. Therefore, the wealth can also be viewed as a weight when resources are allocated.

The competitive market always seeks the equilibrium price that causes supply to equal demand. The equilibrium price can be determined directly; however, in a decentralized economy some terms (the utility and wealth of each agent) may not be known. For this reason, the equilibrium price is determined via a tâtonnement process [103]². First proposed by Léon Walras, the tâtonnement process iteratively adjusts the price with respect to excess demand. The excess demand is a function of the total (aggregate) demand and supply of the resource. The price increases if the demand is greater than the supply and decreases when the demand is less than the supply. It is important to note that the demand and supply at the current price must be known before an adjustment can occur. The iterative process repeats until a price is reached such that supply equals demand; at this point the market and price are in equilibrium. This is referred to as “clearing the market,” where

¹Not to be confused with the indifference curve which is normally convex.

²Alternatively, an auction or bidding procedure can be used [76].

consumers maximize utility given their budget constraints and producers maximize profits. Refer to the prices calculated before the equilibrium price is reached as *intermediate prices*. Buying and selling normally do not occur with the intermediate prices [100]; however, this constraint will not apply to the bandwidth “spot market” (described in the next chapter), since bandwidth in this market is considered a non-storable resource. This allows demands to change dynamically and is achieved using a modified tâtonnement process. Once the market is in equilibrium the resulting allocation is Pareto-optimal and weighted max-min fair, which is proven in the next sections.

5.2 Fairness and Optimality

The allocation provided by an economy consisting of multiple independent competitive markets in equilibrium can be described as *efficient* (Pareto-optimal) and *weighted max-min fair*. Furthermore, with appropriate wealth distribution the allocation is also *equitable*. This section formally defines the terms efficient (Pareto-optimal), weighted max-min fair and equitable allocations. Then three theorems are introduced, that indicate conditions under which an economy consisting of multiple competitive markets can achieve these important goals.

5.2.1 Pareto-Optimal Allocations and Weighted Max-Min Fairness

Assume an economy consists of a set of independent competitive markets L . Each market i sells only a unique type of resource with supply s^i . This implies that an array of prices exists $\{p\}$ in the economy, where the price for the resource sold at market i is p^i . All consumers in the economy belong to set A , where consumer j desires resources belonging to the set $R^j \subset L$. Consumer j has an amount of wealth for each market $i \in R^j$. Denote $w^{j,i}$ as the amount allotted for market i by consumer j ; in addition, assume this amount is equal for all markets the consumer participates in ($w^{j,h} = w^{j,i}$, $\forall h, i \in R^j$). Therefore, the second superscript of $w^{j,i}$ (i , indicating the market) will be dropped for brevity. Each consumer has a maximum amount b^j which is desired for any resource. Let $a^{j,i}$ be the allocation for consumer j in market i . Furthermore, assume the consumer must

purchase the same amount in each market³ ($a^{j,h} = a^{j,i}$, $\forall h, i \in R^j$). As done for w^j , the second superscript of $a^{j,i}$ (indicating the market) will be dropped⁴. Denote A^i as the set of consumers participating in market i . Consumer j is either “completely satiated” or “non-satiated” with their allocation a^j at market i . Let C^i be the set of completely satiated consumers and N^i be the set of non-satiated consumers at market i ; therefore, $C^i \cup N^i = A^i$ and $\sum_{j \in A^i} a^j \leq s^i$ must always be true for all markets in the economy.

Definition 5.2.1. Completely satiated: At market i with price p^i , consumer j is completely satiated with a^j if the amount of resources affordable is greater than what is desired, b^j .

$$\text{if } \frac{w^j}{p^i} \geq b^j \quad \text{then } a^j = b^j \quad (5.1)$$

Definition 5.2.2. Non-satiated: At market i with price p^i , consumer j is non-satiated with a^j if the amount of resources affordable is less than or equal to what is desired, b^j .

$$\text{if } \frac{w^j}{p^i} < b^j \quad \text{then } a^j = \frac{w^j}{p^i} \quad (5.2)$$

Consumer j will purchase resources from each market $i \in R^j$. Depending on the price associated with each market, the consumer can afford different amounts. As previously mentioned, assume the consumer will always purchase the same amount at each market $i \in R^j$. This amount a^j is equal to the minimum amount that is affordable at any market $i \in R^j$ (but no more than the maximum desired b^j),

$$a^j = \min \left\{ \min_{i \in R^j} \left\{ \frac{w^j}{p^i} \right\}, b^j \right\} \quad (5.3)$$

The market in R^j with the highest price is considered *saturated* for consumer j , since only the minimum amount of resources can be purchased (which ultimately determines the amount to purchase at the remaining markets in R^j). At the saturated market the consumer is non-satiated; however, for the remaining markets in R^j the consumer is considered

³This assumption becomes clear when the economy is a computer network and the resource is link bandwidth.

⁴The requirements ($w^{j,h} = w^{j,i}$, $\forall h, i \in R^j$) and ($a^{j,h} = a^{j,i}$, $\forall h, i \in R^j$) can be removed and weighted max-min fair and equitable allocations can be proved for individual markets (instead of an entire economy).

satiated. For example, assume R^j consists of three markets and the consumer can afford 10 units at market 1, 5 units at market 2 and 20 units at market 3. Market 2 is saturated and the consumer will only purchase 5 units at each market. In the case where the consumer can afford b^j at each market in R^j , then the consumer is considered completely satiated at each market in R^j .

Definition 5.2.3. Feasibility: For competitive market i , the price and an allocation array, $[p^i, \{a\}]$, are said to be feasible if and only if,

$$(i) \quad s^i = \sum_{j \in A^i} a^j$$

N.B. The case where $s^i > \sum_{j \in A^i} a^j$ is not considered since resources are not scarce.

$$(ii) \quad p^i \cdot a^j \leq w^j \quad \forall j \in A^i$$

Definition 5.2.4. Competitive equilibrium: At price p_*^i and allocation array $\{a\}$, competitive market i is in equilibrium if and only if,

$$(i) \quad [p_*^i, \{a\}] \text{ is feasible}$$

$$(ii) \quad w^j(a^j) \geq w^j(\hat{a}^j) \text{ for all } \hat{a}^j, \text{ where } \hat{a}^j \leq b^j \text{ and } p_*^i \cdot \hat{a}^j \leq w^j, \text{ such that } p_*^i \cdot a^j \geq p_*^i \cdot \hat{a}^j \text{ for all } j \in A^i$$

Lemma 5.2.1. *If $[p_*^i, \{a\}]$ is the allocation of competitive market i in equilibrium, then the following is true*

$$\frac{a^j}{w^j} = \frac{a^k}{w^k}, \quad \forall j, k \in N^i \quad (5.4)$$

Proof. Assume $[p_*^i, \{a\}]$ is the allocation of competitive market i in equilibrium, and $j, k \in N^i$. From definition 5.2.2, the allocation of non-satiated consumers is,

$$a^j = \frac{w^j}{p_*^i}, \quad a^k = \frac{w^k}{p_*^i} \quad (5.5)$$

From lemma 5.2.1,

$$\frac{\frac{w^j}{p_*^i}}{w^j} = \frac{\frac{w^k}{p_*^i}}{w^k} \quad \rightarrow \quad \frac{1}{p_*^i} = \frac{1}{p_*^i} \quad (5.6)$$

□

Lemma 5.2.2. *If $[p_*^i, \{a\}]$ is the allocation of competitive market i in equilibrium, then the following is true*

$$\max_{\forall j \in C^i} \left\{ \frac{a^j}{w^j} \right\} \leq \frac{a^k}{w^k}, \quad \forall k \in N^i \quad (5.7)$$

Proof. Assume $[p_*^i, \{a\}]$ is the allocation of competitive market i in equilibrium. Denote $a^j = \max_{\forall j \in C^i} \{a^j/w^j\}$ and $k \in N^i$. Suppose contrary to lemma 5.2.2 that, $\frac{a^j}{w^j} > \frac{a^k}{w^k}$. Substituting for a^j and a^k ,

$$\frac{b^j}{w^j} > \frac{w^k}{p_*^i} \rightarrow \frac{b^j}{w^j} > \frac{1}{p_*^i} \quad (5.8)$$

From the definition 5.2.1,

$$b^j \leq \frac{w^j}{p_*^i} \quad (5.9)$$

Dividing both sides by w^j

$$\frac{b^j}{w^j} \leq \frac{w^j}{p_*^i w^j} \rightarrow \frac{b^j}{w^j} \leq \frac{1}{p_*^i} \quad (5.10)$$

Combining equations 5.8 and 5.10,

$$\frac{1}{p_*^i} < \frac{b^j}{w^j} \leq \frac{1}{p_*^i} \quad (5.11)$$

which is not feasible.

□

Definition 5.2.5. Pareto Optimality: The feasible allocation array $\{a\}$ is said to be Pareto-optimal if there does not exist another feasible allocation array $\{\hat{a}\}$, such that $u^j(\hat{a}^j) \geq u^j(a^j) \quad \forall j \in A$ with a strict inequality for at least one j .

Theorem 5.2.3. *The allocation of an economy consisting of independent competitive markets in equilibrium, $[\{p_*\}, \{a\}]$, is Pareto-optimal.*

Proof. A proof that the allocation of an economy consisting of independent competitive markets in equilibrium is Pareto-optimal is given in [100]. This proof can be used for the

economy presented in this chapter (with a slight modification). For completeness, the modified proof is presented here.

Suppose $\{a\}$ is not Pareto-optimal. Then there exists $\{\hat{a}\}$ where

- (i) $[\{p_*\}, \{\hat{a}\}]$ is feasible
- (ii) $u^j(\hat{a}^j) \geq u^j(a^j)$ for all $j \in A$
- (iii) $u^j(\hat{a}^j) > u^j(a^j)$ for at least one j

From definition 5.2.4 (ii) we have

$$p_*^i \cdot \sum_{j \in A^i} \hat{a}^j > p_*^i \cdot \sum_{j \in A^i} a^j \quad (5.12)$$

However, definition 5.2.3, condition (i) requires

$$p_*^i \cdot \sum_{j \in A^i} a^j = p_*^i \cdot s^i \quad (5.13)$$

Therefore we have

$$p_*^i \cdot \sum_{j \in A^i} \hat{a}^j > p_*^i \cdot s^i \quad (5.14)$$

which contradicts the feasibility of $\{\hat{a}\}$. □

Definition 5.2.6. Weighted max-min fair: An allocation of resources $\{a\}$ with weights $\{w\}$ is weighted max-min fair if it is feasible, and if, for any other feasible allocation $\{\hat{a}\}$,

$$\exists j : \hat{a}^j > a^j \implies \exists k : \frac{\hat{a}^k}{w^k} < \frac{a^k}{w^k} \leq \frac{a^j}{w^j} \quad (5.15)$$

Theorem 5.2.4. *The allocation of an economy consisting of independent competitive markets in equilibrium is weighted max-min fair, where the weight of each consumer is their wealth.*

Proof. Assume $[\{p_*\}, \{a\}]$ is the allocation of an economy consisting of independent competitive markets in equilibrium. Let $\{\hat{a}\}$ be any other feasible allocation, where $\hat{a}^j = a^j + \delta^j \geq 0$ and $\sum \delta^j = 0$. Only non-satiated consumers may increase their allocation, requiring other consumers(s) to decrease their allocation. Let two consumers j and k participate in market i ($i \in R^j, R^k$). Assume consumer j is a non-satiated and gains resources under $\{\hat{a}\}$ implying $\delta^j > 0$. Denote consumer k as a consumer that loses resources under $\{\hat{a}\}$ implying $\delta^k < 0$. Consider two cases, (i) consumer k is satiated and, (ii) consumer k is non-satiated.

Case (i), consumer $k \in C^i$.

Combining the assumptions above with lemma 5.2.2

$$\hat{a}^j > a^j \quad \text{and} \quad \frac{\hat{a}^k}{w^k} < \frac{a^k}{w^k} \leq \frac{a^j}{w^j} \quad (5.16)$$

which satisfies the requirement for weighted max-min fairness.

Case (ii), consumer $k \in N^i$.

Combining the assumptions above lemma 5.2.1

$$\hat{a}^j > a^j \quad \text{and} \quad \frac{\hat{a}^k}{w^k} < \frac{a^k}{w^k} = \frac{a^j}{w^j} \quad (5.17)$$

which satisfies the requirement for weighted max-min fairness.

□

To provide perspective to the different types of fairness and optimality, consider all the possible weighted max-min fair allocations as a set. Each member of this set represents the allocation achieved with a certain wealth distribution. Given the conditions required for a weighted max-min allocation and the shape of the utility curve, each member of the set is Pareto-optimal (the conditions required for theorem 5.2.4 include those for theorem 5.2.3). A max-min fair allocation is a member of this set, where the wealth of each consumer is equal. In addition, an equitable allocation (defined in section 5.2.2) is also a member of this set, where the wealth distribution results in equal utility for each consumer.

An Alternative Weighted Max-Min Fair Proof

The fairness proofs introduced in this chapter are based on the competitive market model and are defined in a microeconomic context. However, since the economic model will be used for network resource allocation in chapters 6 and 7, the “network-oriented” fairness proofs described by [6, 45] can be used (with some modifications) to prove weighted max-min fairness. When using these proofs, users are consumers and links are markets.

In [6, 45] the following proposition is made; an allocation $\{a\}$ is max-min fair if every user has a bottleneck link. This proposition depends on the definition of a bottleneck link, which has two parts [6]. First, if any user considers link i a bottleneck then the entire capacity of the link must be allocated $\sum_{j \in A^i} a^j = s^i$. Second, if user j considers link i a bottleneck then $a^j \geq a^k$, $\forall k \in A^i$. However, this proposition and definition does not apply to weighted max-min fairness, and does not permit users to have a maximum desired allocation b^j (as done in the economic model). Therefore, changes must be made to apply the proposition to the economic model.

A modification to the bottleneck definition is required to apply the proposition to weighted max-min fairness [45]. Accounting for the weights (wealths) of each user, the second part of the bottleneck definition becomes; if user j considers link i a bottleneck then $\frac{a^j}{w^j} \geq \frac{a^k}{w^k}$, $\forall k \in A^i$. The first part of the bottleneck definition remains the same. To account for the maximum desired allocation of each user b^j , it is suggested in [6] that a fictitious link be added to the end of the route of each user. Each fictitious link will have capacity equal to b^j , which forces each user to have a bottleneck link (a requirement for the proposition, but not for the microeconomic-based theorem 5.2.4).

Using the modified proposition and the addition of fictitious links, the economic model can be proven to achieve weighted max-min fair allocations. First, lemmas 5.2.1 and 5.2.2 are required to prove the allocation of a competitive market in equilibrium adheres to the bottleneck definition. Due to the fictitious links, every user will have a bottleneck link. For that reason, the proposition can be used to prove the allocation of an economy consisting of multiple competitive markets in equilibrium is weighted max-min fair.

5.2.2 Equitable Allocation

A Pareto-optimal resource allocation in microeconomics is called *efficient*, and many different efficient allocation exist for a competitive market in equilibrium (consider the different possible allocations of wealth) [76]. For this reason, a social welfare criterion, the *equitable* criterion, is used to compare and rank efficient allocations. In economics, the equitable criterion states that each user in the economy should enjoy approximately the same level of utility [76]. This definition must be extended to apply to an economy consisting of multiple independent competitive markets. For such an economy the equitable criterion states that users, who share a common saturated market, must enjoy approximately the same level of utility.

Definition 5.2.7. Equitable allocation: An allocation of resources $\{a\}$ is equitable if it is feasible, and if, for any other feasible allocation $\{\hat{a}\}$,

$$\exists j : u^j(\hat{a}^j) > u^j(a^j) \implies \exists k : u^k(\hat{a}^k) < u^k(a^k) \leq u^j(a^j) \quad (5.18)$$

It is important to note this does not necessarily correspond to equal amounts of a resource (the goal of max-min). In a network economy, this can also be referred to as a “QoS-fair” or “utility-fair” allocation. An equitable allocation can be achieved by a competitive market in equilibrium when the wealth of each consumer is correctly assigned. This is described next⁵.

Consumer j has utility function $u^j(a^j)$ that indicates a utility value q^j for an allocation amount a^j . The inverse of the utility function, denoted as $\bar{u}^j(q^j)$, indicates an allocation amount a^j that achieves a utility value of q^j . Define the aggregate inverse utility function for all consumers who participate in and consider market i saturated as,

$$\bar{v}^i(\cdot) = \sum_{j \in A^i} \bar{u}^j(\cdot) \quad (5.19)$$

Since $\bar{u}^j(\cdot)$ is monotonic, $\bar{v}^i(\cdot)$ is monotonic and has a unique solution for any feasible utility value. At equilibrium the supply equals the demand; let q_*^i be the utility value for

⁵A method for determining weights in a “Fair Queueing” wireless scheduler presented in [7] can be viewed as wealth distribution technique; however, the method would only apply to a single market not an economy consisting of multiple markets.

all consumers at which this occurs, i.e.,

$$s^i = \bar{v}^i(q_*^i) = \sum_{j \in A^i} \bar{u}^j(q_*^i) \quad (5.20)$$

q_*^i can be found quite easily, since $\bar{u}(\cdot)$ is monotonic. To provide each consumer the same utility level q_*^i when the market is in equilibrium, the wealth of consumer j is set as follows:

$$w^j = \bar{u}^j(q_*^i) \quad (5.21)$$

The previous description determined the wealth distribution that achieves an equitable allocation for a single competitive market. Using this as a basis, algorithm 5.1 determines the wealth distribution that achieves an equitable allocation for an entire economy consisting of multiple independent markets.

Algorithm 5.1 requires the utility curve and route of each consumer in the economy. Acquiring such information reliably may not be possible; therefore, the algorithm may not be applicable to an actual network. The algorithm is presented for completeness of this section. Approximations of algorithm 5.1, that require far less information, are presented and used in section 6.4.1.

Lemma 5.2.5. *If $\{w\}$ is the wealth allocation provided by algorithm 5.1 and $\{a\}$ is the allocation of competitive market i in equilibrium, then the following is true,*

$$\max_{\forall j \in C^i} \{u^j(a^j)\} \leq u^k(a^k), \quad \forall k \in N^i \quad (5.22)$$

Proof. Assume an economy consists of two independent competitive markets $L = \{h, i\}$ and two consumers $A = \{j, k\}$. Let consumer j participate in markets $R^j = \{h, i\}$ and consumer k participate in market $R^k = \{i\}$. Furthermore, assume on the first iteration of algorithm 5.1, market h has the lowest utility ($q_*^h < q_*^i$). User j is assigned a wealth that will yield a utility of q_*^h and is a member of sets N^h and C^i . On the second iteration, market i has the lowest utility. User k is assigned a wealth that will a utility of q_*^h and remains a member of set N^i . Therefore, once the markets have reached equilibrium

$$q_*^h < q_*^i \rightarrow u^j(a^j) < u^k(a^k) \quad \text{where } j \in C^i, \quad k \in N^i \quad (5.23)$$

□

Algorithm 5.1 Wealth calculation algorithm for an equitable allocation.

```

1:  /**** variable initialization ****/
2:   $D \leftarrow L$  /* set of markets */
3:  for all  $i \in L$  do
4:     $C^i \leftarrow \emptyset$ 
5:    for all  $j : i \in R^j$  do
6:       $N^i = N^i \cup j$  /* assume all consumers of market  $i$  are non-satiated */
7:    end for
8:  end for
9:  /**** start wealth calculation algorithm ****/
10: while  $D \neq \emptyset$  do
11:    $q_{min} = \infty$ 
12:   for all  $i : i \in D$  do
13:     calculate  $q_*^i$  using consumers in  $N^i$ 
14:     /* determine market with smallest  $q_*^i$  */
15:     if  $q_*^i \leq q_{min}$  then
16:        $q_{min} = q_*^i$ 
17:        $h = i$ 
18:     end if
19:   end for
20:   /* assign wealth to all consumers participating in */
21:   /* and who are non-satiated with market  $h$  */
22:   for all  $j : h \in R^j$  and  $j \in N^h$  do
23:      $w^j = \bar{u}^j(q_*^h)$ 
24:     /* consumer is satiated w.r.t. remaining markets in  $R^j$  */
25:     for all  $i : i \in R^j$  and  $i \neq h$  do
26:        $C^i \leftarrow C^i \cup j$ 
27:        $N^i \leftarrow N^i \setminus j$ 
28:     end for
29:   end for
30:    $D \leftarrow D \setminus h$  /* market  $h$  has been processed, remove from set */
31: end while

```

Theorem 5.2.6. *Allocating wealth using algorithm 5.1 yields an equitable allocation for an economy consisting of independent competitive markets in equilibrium.*

Proof. Assume $[\{p_*\}, \{a\}]$ is the allocation of an economy consisting of independent competitive markets in equilibrium, where the wealth of consumers $\{w\}$ was determined from algorithm 5.1. Let $\{\hat{a}\}$ be any other feasible allocation, where $\hat{a}^j = a^j + \delta^j \geq 0$ and $\sum \delta^j = 0$. Only non-satiated consumers may increase their allocation, requiring other consumer(s) to decrease their allocation. Let two consumers j and k participate in market i ($i \in R^j, R^k$). Assume consumer j is a non-satiated (considers market i saturated) and gains resources under $\{\hat{a}\}$ implying, $\delta^j > 0$ and $u^j(\hat{a}^j) > u^j(a^j)$. Denote consumer k as a consumer that loses resources under $\{\hat{a}\}$ implying, $\delta^k < 0$ and $u^k(\hat{a}^k) < u^k(a^k)$. Consider two cases, (i) consumer k is satiated and, (ii) consumer k is non-satiated.

Case (i), consumer $k \in C^i$.

Combining the assumptions above with lemma 5.2.5

$$u^j(\hat{a}^j) > u^j(a^j) \quad \text{and} \quad u^k(\hat{a}^k) < u^k(a^k) \leq u^j(a^j) \quad (5.24)$$

which satisfies the requirement for an equitable allocation.

Case (ii), consumer $k \in N^i$.

As specified in algorithm 5.1, all non-satiated consumers of market i receive the same utility.

Combining this with the assumptions above

$$u^j(\hat{a}^j) > u^j(a^j) \quad \text{and} \quad u^k(\hat{a}^k) < u^k(a^k) = u^j(a^j) \quad (5.25)$$

which satisfies the requirement for an equitable allocation.

□




Figure 5.1: Network configuration for the fairness examples.

5.3 Example Competitive Market Allocations

In this section, examples of weighted max-min and equitable allocations are given for a simple economy consisting of two markets. For each example assume the economy is the network given in figure 5.1. This network consists of three users, two switches and two links, where each link has a total capacity of ten units. Users 0 and 1 use links 0 and 1 (in that order), while user 2 uses only link 1. Users are considered consumers in the economy and the switches are the producers. Switches sell link bandwidth to the users; therefore, switch 0 sells link 0 bandwidth and switch 1 sells link 1 bandwidth. Each link is considered an independent competitive market, which is the economic model described in section 5.2.

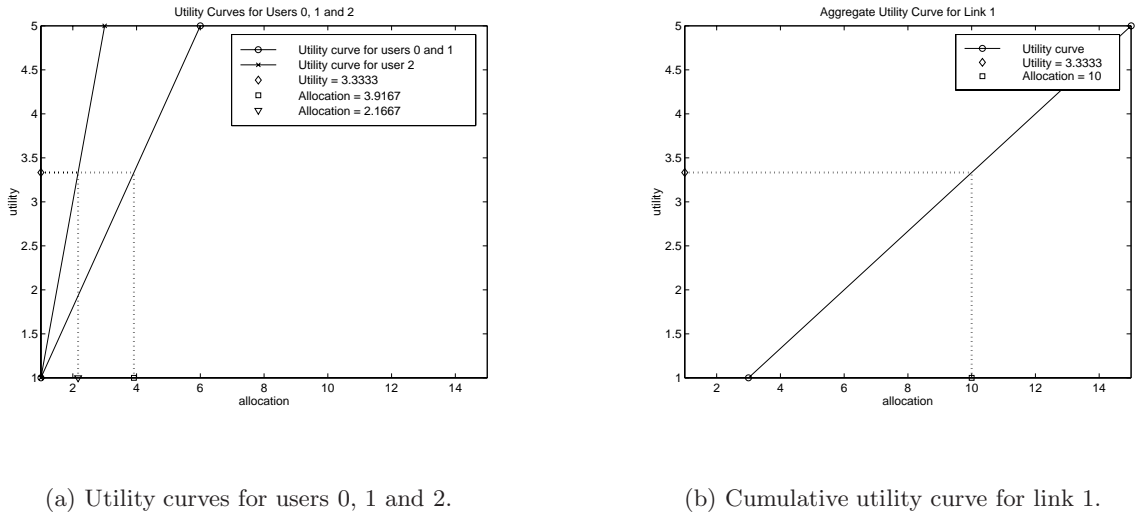
5.3.1 Weighted Max-Min Fair

Assume the users have the wealths and maximum demands given in table 5.1. Assuming the markets have reached equilibrium, the equilibrium price⁶ for link 0 bandwidth is $\frac{3}{5}$ and the equilibrium price for link 1 bandwidth is $\frac{2}{3}$. As defined in section 5.2.1, all users find that link 1 is their “saturated” market. At this link, user 2 can afford 3 units of bandwidth; however the maximum demand is 1 unit. For this reason, user 2 is considered “completely satiated” according to definition 5.2.1. Users 0 and 1 can only afford 6 and 3 units respectively at link 1; therefore, these users are considered “non-satiated” as defined by definition 5.2.2. The final allocations are given in table 5.1 and are weighted max-min fair as defined by 5.2.6.

⁶How the equilibrium price is determined is given in chapter 6

User	Wealth w^j	Demand b^j	Allocated a^j
0	4	10	6
1	2	10	3
2	2	1	1

Table 5.1: Example weighted max-min fair allocation.



(a) Utility curves for users 0, 1 and 2.

(b) Cumulative utility curve for link 1.

Figure 5.2: Utility curves for the equitable allocation example.

5.3.2 Equitable Allocations

As discussed in section 5.2.2, an equitable allocation measures fairness in terms of the utility obtained from the resources. To obtain an equitable allocation, the wealth of each user must be distributed according to algorithm 5.1. This algorithm requires the route and utility curve of each consumer (user) in the economy. As described in section 5.2.2, utility curves are assumed to be continuous and monotonic, as seen in figure 5.2. The horizontal axis of the utility curve measures utility (satisfaction) as a real number, while the horizontal axis measures the corresponding allocation amount. For example, users 0 and 1 require 6 units of bandwidth to receive the highest possible utility (5).

To determine the equitable allocation, the utility q_*^i (equation 5.20) must be calculated for the two links (step 12 of algorithm 5.1). For link 0 q_*^0 is 4.2, and for link 1 q_*^1

User	Demand b^j	Wealth w^j	Allocated a^j	Utility u^j
0	6	3.9167	3.9167	3.3333
1	6	3.9167	3.9167	3.3333
2	3	2.1667	2.1667	3.3333

Table 5.2: Example equitable allocation.

is 3.3333. Therefore, link 1 is the saturated market and all users of this link will have a utility of 3.3333. The wealth of each user must result in a utility of 3.3333. As seen in figure 5.2(a), the wealth for users 0 and 1 should be 3.9167 (resulting in a QoS score of 3.3333). Similarly, the wealth for user 2 should be 2.1667. Note that the wealth distribution equals the final allocation for each user. This is expected since one unit of currency is exchanged for one unit of bandwidth.

5.4 Chapter Summary

In this chapter, the competitive market model was discussed. This model consists of two types of agents, consumers and producers. Consumers purchase resources to maximize utility (happiness), while producers sell or rent resource at the market price to maximize profits. When a price is determined that causes supply to equal demand; the market and price are in equilibrium. At equilibrium, consumers maximize utility given their budget constraints and producers maximize profits.

This chapter also discussed a model that consists of multiple independent competitive markets. When the markets are in equilibrium, it was proven that optimal and fair allocations are obtained. Possible fairness measures include, weighted max-min fair and equitable allocations. In addition, a method of wealth distribution is described that achieves an equitable allocation.

The competitive market model is used as the basis for the multi-user allocation methods described in chapters 6 and 7. This model provides a simple method for allocating network resources, that achieves efficient and fair allocations.