Pareto Efficiency
(also called Pareto Optimality)

1 Definitions and notation

Recall some of our definitions and notation for preference orderings. Let X be a set (the set of *alternatives*); we have the following definitions:

1. A **relation** R on X is a subset of $X \times X$. We often write xRy instead of $(x, y) \in R$ and we say "x is R-related to y".

2. If R is a relation on X, we denote the **complement** of R by $\not R$ (instead of $\sim R$, because $\sim$ will be given another meaning). Thus, $x \not R y$ means that x is *not* R-related to y: $(x, y) \notin R$.

3. A **strict preordering** of X is a transitive and irreflexive relation P on X. We usually write $x \succ y$ for xPy, and also $y \prec x$. We say that x is "preferred" to y. A **weak preordering** of X is a transitive and reflexive relation on X, and a **complete preordering** of X is a transitive and complete relation on X. For either one, we usually write $x \succsim y$ and say that x is weakly preferred to y.

4. If P is a strict preordering, we denote the corresponding **indifference relation** by I, defined by $xIy \iff [x \not P y \;\&\; y \not P x]$. We also write $x \sim y$ for xIy, and $x \succsim y$ (also $y \precsim x$) for $[xPy$ or $xIy]$. Note that $\sim$ is both reflexive and symmetric, but it need not be transitive; and that $\succsim$ is complete, but it need not be transitive. (Can you provide a counterexample to show that transitivity may fail?) If $\sim$ *is* transitive , then $\succsim$ is a complete preordering.

5. If $\succsim$ is a complete preordering, then $\sim$ is transitive, and $[x \succsim y \;\&\; y \succ z]$ implies $x \succ z$ for any $x, y, z \in X$.

2 Aggregation of rankings into a single ranking

Let X be a set of **alternatives**, generically denoted by x; let N be a set of n **individuals**, generically denoted by i; and let $\mathcal{P}$ be a set of admissible **preorderings** ("rankings," or "preferences") over X, generically denoted by P. For any list $\mathbf{P} = (P_1, ..., P_n) \in \mathcal{P}^n$ of individual rankings we would like to have a single P that "summarizes" or represents $\mathbf{P}$ — for example, a single P that could be used as the criterion for making decisions that take account of $P_1, ..., P_n$. What we want, then, is an "aggregation rule" or function

$$a : \mathcal{P}^n \longrightarrow \mathcal{P}, \qquad i.e., \quad (P_1, ..., P_n) \overset{a}{\longmapsto} \overline{\mathbf{P}}. \tag{1}$$

We want to have a rule we can use to aggregate a list $\mathbf{P} = (P_1, ..., P_n)$ of individual rankings into a single "aggregate" ranking, $\overline{\mathbf{P}}$. In other words, we want to have an aggregation function or rule

$$a : \mathcal{P}^n \longrightarrow \mathcal{P}, \qquad i.e., \quad (P_1, ..., P_n) \overset{a}{\longmapsto} \overline{\mathbf{P}}. \tag{2}$$

(Note the similarity with the notation for the sample mean of a list of n numbers: $\overline{x} = \frac{1}{n}\sum_i x_i$. The sample mean is a way of aggregating a list of numbers into a single "representative" number — *i.e.*, it's a function that maps a list of numbers into a single number).

Instead of framing the problem as one of aggregating a list of rankings into a single ranking, we could alternatively frame the problem as one of aggregating a list of utility functions into a single utility function. We will return to this idea in Section 7.

2.1 Examples

Here are several examples of sets X of alternatives for which we might wish to aggregate a list of individual rankings into a "representative" ranking:

1. X is a set of allocations $\mathbf{x} = (\mathbf{x}^1, ..., \mathbf{x}^n) \in \mathbb{R}^{nl}_+$.
2. X is a set of candidates for a job, or for a political position.
3. X is a set of public policies.
4. X is a set of teams; for example, $X = \{A, B, C\}$, A : Arizona, B : Boston College, C : California.
5. X is a set of tennis players; for example, $X = \{A, B, C\}$, A : Agassi, B : Becker, C : Chang.

In this last example the n individual rankings $P_1, ..., P_n$ could be the rankings (*i.e.*, the order of finish) in each of n tournaments, and the problem is to aggregate these tournament results into a single ranking of the players. This is exactly what the "ATP Ranking" is, an aggregation of the players' tournament finishes during the preceding year into a single ranking of the players. The ATP ranking uses a specific rule (function, algorithm) $a : (P_1, ..., P_n) \longrightarrow \bar{P}$ to calculate $\bar{P}$. The ATP rule weights the various tournaments differently, assigning more weight for example to the so-called Grand Slam tournaments than to other tournaments. (ATP is the abbreviation used by the Association of Tennis Professionals.)

3 The Pareto ranking

Definition: Let $\mathbf{P} = (P_1, ..., P_n)$ be a list of preorderings of a set X of alternatives. We say that $\widetilde{x}$ is a **Pareto improvement** upon x (which we write $\widetilde{x}\overline{\mathbf{P}}x$), or that $\widetilde{x}$ **Pareto dominates** x, if $\forall i : x \not P_i \widetilde{x}$ and $\exists i : \widetilde{x} P_i x$. $\overline{\mathbf{P}}$ is called the **Pareto ranking** or **Pareto ordering** associated with the list $\mathbf{P} = (P_1, ..., P_n)$

Remark: As above, the alternatives (the elements of X) needn't be allocations — they could be political parties, candidates, athletic teams, etc. — and the function $(P_1, ..., P_n) \longmapsto \overline{\mathbf{P}}$ is only one of many possible ways to aggregate the list $(P_1, ..., P_n)$ of rankings into a single "aggregate ranking" $\overline{\mathbf{P}}$.

Remark: If each P_i is transitive or irreflexive, then so is $\overline{\mathbf{P}}$. But even if each I_i is transitive, and each P_i is transitive and irreflexive, $\bar{\mathbf{I}}$ may fail to be transitive, as Examples 2 and 3 below demonstrate, or $\bar{\mathbf{I}}$ may be transitive but uninformative, as Example 1 demonstrates.

In Examples 1 and 2, below, the set of alternatives is $X = \{A, B, C\}$. Here are two possible interpretations of the examples:

> A is Arizona, B is Boston College, C is California. Each P_i could be an individual's ranking of these universities' basketball teams, or their economics departments, or their reputations as party schools, etc.
>
> A is Agassi, B is Becker, C is Chang. Each P_i is their order of finish in a tournament.

Example 1: $X = \{A, B, C\}; A \succ_1 B \succ_1 C; B \succ_2 C \succ_2 A; C \succ_3 A \succ_3 B$. Therefore $A\bar{\mathbf{I}}B$, $B\bar{\mathbf{I}}C$, and $C\bar{\mathbf{I}}A$. Thus, the aggregate indifference relation $\bar{\mathbf{I}}$ is transitive, but not very useful.

Example 2: $X = \{A, B, C\}; A \succ_1 C \succ_1 B; B \succ_2 A \succ_2 C$. Therefore $A \sim B, B \sim C$, and $A \succ C$. Thus, the aggregate indifference relation $\bar{\mathbf{I}}$ (or $\sim$) is not transitive.

Example 3: Figure 1 is an Edgeworth box with two consumers, each with a standard preference. Nevertheless, the rankings of the allocations A, B, and C are just as in Example 2: $A \succ_1 C \succ_1 B$ and $B \succ_2 A \succ_2 C$, so that the Pareto ranking is also the same as in Example 2: $A \sim B, B \sim C$, and $A \succ C$.

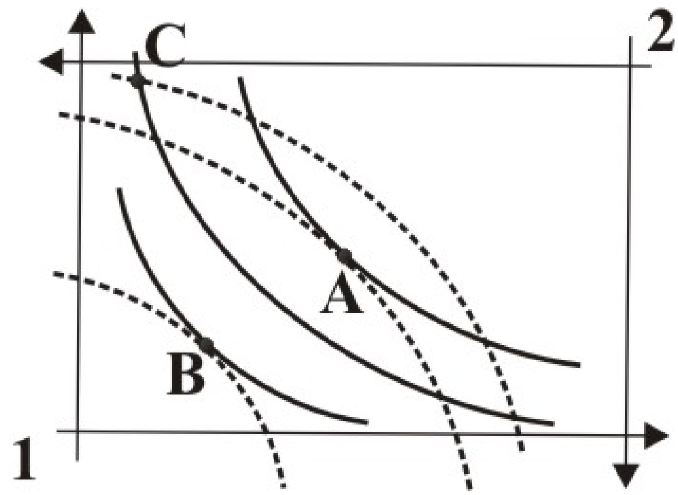

4 Pareto Efficiency

Definition: Let X be a set of alternatives, and let $(\succsim_i)_1^n$ be a list of preferences over X. An alternative $\widehat{x}$ is **Pareto efficient** if no alternative in X Pareto dominates $\widehat{x}$.

We often have a "fundamental" set X of alternatives — for example, all the conceivable or well defined alternatives — but only a subset $\mathcal{F} \subseteqq X$ of the alternatives are actually possible, or **feasible**. Moreover, we generally want to allow the set $\mathcal{F}$ to vary and to see how the Pareto efficient alternatives depend on the set $\mathcal{F}$. Our standard allocation problem is a good example of this: we take $X = \mathbb{R}_+^{nl}$ to be the set of all conceivable allocations, and this is the set over which individuals' preferences $\succsim_i$ are defined; but in order to say whether a given allocation $(\mathbf{x}^i)_1^n$ is efficient we don't want to insist that it not be dominated by any *conceivable* allocation, only that it not be dominated by any other *feasible* allocation — *i.e.*, by any other allocation that can actually be achieved with existing resources.

Definition: Let $(\succsim_i)_1^n$ be a list of preferences over a set X, and let $\mathcal{F} \subseteqq X$. An alternative $x \in \mathcal{F}$ is **Pareto efficient** (with respect to $\mathcal{F}$) if it is not Pareto dominated by any other alternative $\widetilde{x} \in \mathcal{F}$.

5 Characterizing Pareto efficient allocations

The definition of Pareto efficiency is pretty awkward and clumsy to work with analytically. We'd like to be able to *characterize* the efficient alternatives in some way that's more analytically tractable or more economically intuitive — for example, as the solution to an optimization problem, or in terms of marginal rates of substitution. For our economic allocation problem we can actually establish such a characterization.

So far, with the exception of Example 3 above, we've been dealing with alternatives in the abstract: the alternatives could be just about anything. Defined at this level of generality, the idea of Pareto efficiency can be applied in many useful contexts. But for the economic allocation problem we're studying, the alternatives we want to compare are alternative *allocations*; moveover, when we're dealing with allocations, the individual preferences are typically representable by utility functions. With the structure provided by Euclidean space (where the allocations live) and by using functions instead of orderings, it's pretty easy to characterize the Pareto efficient allocations as the solutions to a constrained maximization problem. We'll tackle that first. And then, since we already know how to characterize the solutions of a constrained maximization problem in terms of first-order conditions, we'll have solved the problem of characterizing the Pareto efficient allocations (or simply the **Pareto allocations**) by first-order conditions. Then we'll find that it's pretty straightforward to translate the first-order conditions into a set of economic marginal conditions — thus giving us a characterization of the Pareto allocations in terms of marginal conditions.

Let's begin by taking an allocation $\widehat{\mathbf{x}} = (\widehat{\mathbf{x}}^i)_1^n$ that's Pareto efficient and we'll show that because $\widehat{\mathbf{x}}$ is a Pareto allocation it must be a solution to a specific constrained maximization problem. The constraints are of course the usual resource constraints, to which we add the requirement that any alternative allocation $\mathbf{x}$ must make $n-1$ of the consumers no worse off than they would have been at $\widehat{\mathbf{x}}$. Then, since $\widehat{\mathbf{x}}$ is Pareto efficient, it must be providing the remaining consumer with the greatest utility possible among all these alternatives $\mathbf{x}$. Note that since $\widehat{\mathbf{x}}$ is given, each $u^i(\widehat{\mathbf{x}}^i)$ in the following proposition is just a real number. Throughout this section we're assuming that the preferences are representable by utility functions.

Proposition: If the allocation $\widehat{\mathbf{x}}$ is Pareto efficient for the endowment bundle $\mathring{\mathbf{x}} \in \mathbb{R}^l_{++}$ and the utility functions $u^1, \dots u^n$ then $\widehat{\mathbf{x}}$ is a solution of the following maximization problem:

$$\begin{aligned} &\max_{(x_k^i)\in\mathbb{R}^{nl}_+} u^1(\mathbf{x}^1) \\ &\text{subject to} \quad x_k^i \geqq 0, \quad i=1,...,n, \quad k=1,...,l \\ &\qquad \sum_{i=1}^n x_k^i \leqq \mathring{x}_k, \quad k=1,...,l \\ &\qquad u^i(\mathbf{x}^i) \geqq u^i(\widehat{\mathbf{x}}^i), \quad i=2,...,n. \end{aligned} \tag{P-Max}$$

Proof: Suppose $(\widehat{\mathbf{x}}^i)_1^n$ is *not* a solution of (P-Max) — *i.e.*, there exist $\widetilde{\mathbf{x}}^1, ..., \widetilde{\mathbf{x}}^n \in \mathbb{R}^l_+$ for which

$$\begin{aligned} \sum_{i=1}^n \widetilde{x}_k^i &\leqq \mathring{x}_k, \qquad k=1,...,l \\ u^i(\widetilde{\mathbf{x}}^i) &\geqq u^i(\widehat{\mathbf{x}}^i) \quad i=2,...,n \\ u^1(\widetilde{\mathbf{x}}^1) &> u^1(\widehat{\mathbf{x}}^1). \end{aligned}$$

Then $(\widetilde{\mathbf{x}}^i)_1^n$ is clearly a Pareto improvement on $(\widehat{\mathbf{x}}^i)_1^n$; *i.e.*, $(\widehat{\mathbf{x}}^i)_1^n$ is not Pareto efficient. ■

A striking feature of this proposition is that it requires no assumptions about the consumers' utility functions. They needn't be convex, or continuous, or even increasing. And the same proof can be used even if the utility functions aren't selfish — *i.e.*, even if some of the consumers care about others' consumption levels.

In order to have a *characterization* of the Pareto efficient allocations, we have to establish the converse of the proposition we've just established: we have to show that any solution of (P-Max) is Pareto efficient. In fact, the converse isn't actually true under such general conditions. But if the consumers' utility functions are all continuous and strictly increasing, that's enough to ensure that the converse is true.

Proposition: If every u^i is continuous and strictly increasing, and if the allocation $\widehat{\mathbf{x}}$ is a solution of the problem (P-Max), then $\widehat{\mathbf{x}}$ is Pareto efficient for $\mathring{\mathbf{x}}$ and the utility functions $u^1, \ldots u^n$.

Proof: Suppose $(\widehat{\mathbf{x}}^i)_1^n$ is *not* Pareto efficient; we will show that then $(\widehat{\mathbf{x}}^i)_1^n$ is not a solution of (P-Max). Since $(\widehat{\mathbf{x}}^i)_1^n$ is not Pareto efficient, there exists a Pareto improvement upon $(\widehat{\mathbf{x}}^i)_1^n$, let's say $(\widetilde{\mathbf{x}}^i)_1^n$:

$$\begin{aligned} \sum_{i=1}^n \widetilde{x}_k^i &\leqq \mathring{x}_k, & k = 1, ..., l \\ u^i(\widetilde{\mathbf{x}}^i) &\geqq u^i(\widehat{\mathbf{x}}^i) & i = 1, ..., n \\ u^i(\widetilde{\mathbf{x}}^i) &> u^i(\widehat{\mathbf{x}}^i) & \text{for some } i. \end{aligned}$$

If $u^1(\widetilde{\mathbf{x}}^1) > u^1(\widehat{\mathbf{x}}^1)$, then $(\widehat{\mathbf{x}}^i)_1^n$ is not a solution of (P-Max), and the proof is complete. So assume that $u^1(\widetilde{\mathbf{x}}^1) = u^1(\widehat{\mathbf{x}}^1)$ and (wlog) $u^2(\widetilde{\mathbf{x}}^2) > u^2(\widehat{\mathbf{x}}^2)$.

Since u^2 is continuous, we may choose $\epsilon > 0$ small enough that every $\mathbf{x}^2 \in B_\epsilon(\widetilde{\mathbf{x}}^2) \cap \mathbb{R}^l_+$ satisfies $u^2(\mathbf{x}^2) > u^2(\widehat{x}^2)$. And since $u^2(\widetilde{\mathbf{x}}^2) > u^2(\widehat{\mathbf{x}}^2)$ and u^2 is strictly increasing, there is some K for which $\widetilde{\mathbf{x}}^2_K > \widehat{\mathbf{x}}^2_K$. Define a new bundle $\overline{\mathbf{x}}^2$ by $\overline{\mathbf{x}}^2_K = \widetilde{\mathbf{x}}^2_K - \frac{1}{2}\epsilon$ and $\overline{\mathbf{x}}^2_k = \widetilde{\mathbf{x}}^2_k$ for $k \neq K$. Then $\overline{x}^2 \in B_\epsilon(\widetilde{\mathbf{x}}^2) \cap \mathbb{R}^l_+$, so $u^2(\overline{\mathbf{x}}^2) > u^2(\widehat{\mathbf{x}}^2)$. Define a new bundle $\overline{\mathbf{x}}^1$ by $\overline{\mathbf{x}}^1 = \widetilde{\mathbf{x}}^1 + \widetilde{\mathbf{x}}^2 - \overline{\mathbf{x}}^2$. And for $i = 3, \ldots, n$, let $\overline{\mathbf{x}}^i = \widetilde{\mathbf{x}}^i$.

Now we have $\overline{\mathbf{x}}^1 + \overline{\mathbf{x}}^2 = (\widetilde{\mathbf{x}}^1 + \widetilde{\mathbf{x}}^2 - \overline{\mathbf{x}}^2) + \overline{\mathbf{x}}^2 = \widetilde{\mathbf{x}}^1 + \widetilde{\mathbf{x}}^2$, so that $\sum_{i=1}^n \overline{\mathbf{x}}^i = \sum_{i=1}^n \widetilde{\mathbf{x}}^i \leqq \mathring{\mathbf{x}}$. We also have $u^i(\overline{\mathbf{x}}^i) \geqq u^i(\widehat{\mathbf{x}}^i), i = 2, ..., n$ and $u^1(\overline{\mathbf{x}}^1) > u^1(\widehat{\mathbf{x}}^1)$. Therefore $(\widehat{\mathbf{x}}^i)_1^n$ is not a solution of (P-max). ■

Combining the two propositions we've just established gives us the following theorem.

Theorem: If every u^i is continuous and strictly increasing, then an allocation $\widehat{\mathbf{x}}$ is Pareto efficient for $\mathring{\mathbf{x}}$ and the utility functions $u^1, \ldots u^n$ if and only if it is a solution of the problem (P-max).

It's important to note that while the problem (P-max) as well as the theorem and the two propositions are all stated in terms of maximizing u^1, the theorem and the propositions are actually true if we restate (P-max) using any one of the n utility functions u^i as the maximand and of course use the remaining $n-1$ utility functions in the constraints. We can see this in either of two ways: each proof can obviously be altered in accordance with the change in the statement of the maximization problem; or we could simply re-index the n individuals in the economy so that the utility function to be maximized becomes u^1, and then the original maximization problem becomes the relevant one.

For interior allocations, we can weaken the requirement that utility functions be strictly increasing, requiring only that they be **locally nonsatiated**.

Definition: A preference $\succeq$ on a set $X \subseteqq \mathbb{R}^l$ is **locally nonsatiated** (LNS) if for any $\mathbf{x} \in X$ and any neighborhood $\mathcal{N}$ of $\mathbf{x}$, there is an $\widetilde{\mathbf{x}} \in \mathcal{N}$ that satisfies $\widetilde{\mathbf{x}} \succ \mathbf{x}$.

Note: We would therefore say that a utility function u on a set $X \subseteqq \mathbb{R}^l$ is locally nonsatiated if for any $\mathbf{x} \in X$ and any neighborhood $\mathcal{N}$ of $\mathbf{x}$, there is an $\widetilde{\mathbf{x}} \in \mathcal{N}$ that satisfies $u(\widetilde{\mathbf{x}}) > u(\mathbf{x})$.

Theorem: If every u^i is continuous and locally nonsatiated, then an interior allocation $\widehat{\mathbf{x}}$ is Pareto efficient for $\mathring{\mathbf{x}}$ and the utility functions $u^1, \ldots u^n$ if and only if it is a solution of the problem (P-max).

Proof: We alter the proof given above, for strictly increasing utility functions, by choosing ϵ small enough that, in addition to having $B_\epsilon(\widetilde{\mathbf{x}}^2)$ in the strict upper-contour set of $\widehat{\mathbf{x}}^2$ we also have both $B_\epsilon(\widetilde{\mathbf{x}}^1) \subset \mathbb{R}^l_+$ and $B_\epsilon(\widetilde{\mathbf{x}}^2) \subset \mathbb{R}^l_+$. Because u^1 is LNS there is a bundle $\overline{\mathbf{x}}^1 \in B_\epsilon(\widetilde{\mathbf{x}}^1) \cap \mathbb{R}^l_+$ for which $u^1(\overline{\mathbf{x}}^1) > u^1(\widehat{\mathbf{x}}^1)$. Define $\overline{\mathbf{x}}^2$ by $\overline{\mathbf{x}}^2 = \widetilde{\mathbf{x}}^2 - (\overline{\mathbf{x}}^1 - \widetilde{\mathbf{x}}^1)$. Then the remainder of the proof proceeds as in the proof for the strictly increasing case. ■

Exercise: Provide a counterexample to show why, for interior allocations, this theorem requires that utility functions be locally nonsatiated, and also a counterexample to show why, at a boundary allocation, local nonsatiation is not enough.

6 Calculus characterization of Pareto efficiency: marginal conditions

Now that we've characterized Pareto efficient allocations as solutions to a constrained maximization problem, it should be straightforward to use that maximization problem to characterize the Pareto allocations in terms of first-order conditions, and then to re-cast the first-order conditions as economic marginal conditions. First-order conditions are calculus conditions, and they require some convexity — *i.e.*, second-order conditions — so throughout this section we assume that each consumer's utility function u_i is continuously differentiable and quasiconcave. To simplify notation we write u^i_k for the partial derivative $\frac{\partial u^i}{\partial x^i_k}$. We also assume that each u^i is strictly increasing:

$u^i_k(\mathbf{x}^i) > 0$ for all i and k. Thus, only those allocations that fully allocate all the goods — those that satisfy $\sum_1^n \mathbf{x}^i = \sum_1^n \mathring{\mathbf{x}}^i$ — could be Pareto allocations. You should be able to verify that under these assumptions the Kuhn-Tucker Theorem's second-order conditions and constraint qualification are satisfied, so that the KT first-order conditions are necessary and sufficient for an allocation $(\widehat{\mathbf{x}}^i)_1^n$ to be a solution of (P-max).

6.1 Interior Allocations

In the previous section we established that an allocation is Pareto efficient if and only if it is a solution of the constrained maximization problem (P-max). Let's assign Lagrange multipliers $\sigma_1, ..., \sigma_l$ to the l resource constraints in problem (P-max) and multipliers $\lambda_2, ..., \lambda_n$ to the $n-1$ utility-level constraints $u^i(\mathbf{x}^i) \geqq u^i(\widehat{\mathbf{x}}^i), i = 2, ..., n$. If all the $\widehat{x}^i_k$'s are strictly positive — *i.e.*, if $(\widehat{\mathbf{x}}^i)_1^n$ is an "interior allocation" — then the first-order marginal conditions for $(\widehat{\mathbf{x}}^i)_1^n$ to be a solution of (P-max) are all equations:

$\exists \lambda_2, ..., \lambda_n \geqq 0$ and $\sigma_1, ..., \sigma_l \geqq 0$ such that for each $k = 1, ..., l$:

$$u^1_k = \sigma_k \quad \text{and} \quad 0 = \sigma_k - \lambda_i u^i_k, \quad i = 2, ..., n \qquad \text{(FOMC)}$$

We can rewrite the last $n-1$ equations as $\lambda_i u^i_k = \sigma_k (i = 2, ..., n; k = 1, ..., l)$. We also have $\sigma_k > 0$ for each k and $\lambda_i > 0$ for each $i = 2, ..., n$ (you should be able to show why this is so; recall that the value of a constraint's Lagrange multiplier is the constraint's shadow value — the marginal increase in the objective value achievable by a one-unit relaxation of the constraint's right-hand-side). Therefore, for every consumer i and every pair of goods k and k', we have

$$\frac{u^i_k}{u^i_{k'}} = \frac{\sigma_k}{\sigma_{k'}}, \quad i.e. \quad MRS^i_{kk'} = \frac{\sigma_k}{\sigma_{k'}}.$$

That last equation says that each consumer's $MRS_{kk'}$ between any two goods k and k' is equal to the relative shadow values of those two goods in the maximization problem (P-max). Clearly then, for any pair of goods every consumer must have the same MRS:

$$MRS^1_{kk'} = ... = MRS^i_{kk'} = ... = MRS^n_{kk'}. \qquad \text{(EqualMRS)}$$

We've derived the equality of MRS's in (EqualMRS) from the Kuhn-Tucker first-order conditions for $(\widehat{\mathbf{x}}^i)_1^n$ to be a solution of (P-max). Therefore (EqualMRS) is a **necessary condition** for $(\widehat{\mathbf{x}}^i)_1^n$ to be a Pareto allocation.

In order to show that (EqualMRS) is also a **sufficient condition** for Pareto efficiency we need to determine values of the Lagrange multipliers σ_k and λ_i for which the equations (FOMC) all hold when the derivatives u_k^i are evaluated at $\widehat{\mathbf{x}}$. Thus,

$$\text{for each } k, \text{ let } \sigma_k = u_k^1(\widehat{\mathbf{x}}^1), \qquad \text{and for each } i, \text{ let } \lambda_i = \frac{\sigma_l}{u_l^i(\widehat{\mathbf{x}}^i)}.$$

For each k and each i we have $\sigma_k > 0$ and $\lambda_i > 0$ and therefore, since the equations (EqualMRS) are satisfied at $(\widehat{\mathbf{x}}^i)_1^n$, we have

$$\frac{u_k^i}{u_l^i} = \frac{u_k^1}{u_l^1} = \frac{\sigma_k}{\sigma_l},$$

which yields

$$\sigma_k = \frac{\sigma_l}{u_l^i} u_k^i = \lambda_i u_k^i,$$

which are exactly the first-order marginal conditions (FOMC) for $(\widehat{\mathbf{x}}^i)_1^n$ to be a solution of (P-max).

We have succeeded in characterizing the interior solutions of (P-max) as the allocations that satisfy the condition (EqualMRS). In the preceding section we characterized the interior Pareto allocations as the solutions to (P-max). Therefore we have the following characterization of the Pareto allocations in terms of marginal conditions:

Theorem: If every u^i is strictly increasing, quasiconcave, and differentiable, then an interior allocation $\widehat{\mathbf{x}}$ is Pareto efficient for $\mathring{\mathbf{x}}$ and the utility functions $u^1, \ldots u^n$ if and only if it satisfies (EqualMRS) and $\sum_1^n \widehat{\mathbf{x}}^i = \sum_1^n \mathring{\mathbf{x}}^i$.

6.2 Boundary Allocations

Typically many of the Pareto allocations are boundary allocations: some consumers' bundles don't include positive amounts of all the goods. We want our marginal conditions to tell us which boundary allocations are Pareto efficient and which aren't, in the same way as the conditions we've just developed do for interior allocations. Since we're dealing with continuous and strictly increasing utility functions, we know that a boundary allocation, just like an interior allocation, is Pareto efficient if and only if it's a solution of (P-max). So all we need to do is adapt the first-order conditions (FOMC) to cover boundary allocations: we have to allow for the equations in (FOMC) to be inequalities when they're associated with variables that have the value zero. Thus, we have

$\exists \lambda_1, ..., \lambda_n \geqq 0$ and $\sigma_1, ..., \sigma_l \geqq 0$ such that for each $k = 1, ..., l$ and each $i = 1, ..., n$:

$$\lambda_i u_k^i \leqq \sigma_k, \quad \text{and} \quad \lambda_i u_k^i = \sigma_k \text{ if } x_k^i > 0 \qquad \text{(FOMC)}$$

Of course these inequalities don't yield the nice equality of all consumers' MRS's for any pair of goods that we obtained in (EqualMRS) for interior allocations. Let's see how these first-order inequalities translate into marginal conditions, for any pair of goods and for any pair of consumers.

Without loss of generality, we consider the two goods $k = 1, 2$. For each consumer (and omitting superscripts for the moment), (FOMC) yields

$$\text{If } x_1 > 0, \text{ then } \frac{u_1}{u_2} \geqq \frac{\sigma_1}{\sigma_2}; \qquad i.e., \, MRS \geqq \frac{\sigma_1}{\sigma_2}. \tag{3}$$

$$\text{If } x_2 > 0, \text{ then } \frac{u_1}{u_2} \leqq \frac{\sigma_1}{\sigma_2}; \qquad i.e., \, MRS \leqq \frac{\sigma_1}{\sigma_2}. \tag{4}$$

Combining (2) and (3) for any two consumers (wlog, let's say they're $i = 1, 2$), we have the following two MRS conditions that must be satisfied at a Pareto efficient allocation:

(A) If $x_1^1 > 0$ and $x_2^2 > 0$, then $MRS^1 \geqq MRS^2$.

(B) If $x_2^1 > 0$ and $x_1^2 > 0$, then $MRS^1 \leqq MRS^2$.

Together, these two conditions cover every combination of positive and zero values of these two goods in the bundles assigned to consumers $i = 1, 2$, as the following table describes. Note that all interior allocations are Case (1) in the table — *i.e.*, the case in which both (A) and (B) above apply, so that we have $MRS^1 = MRS^2$. All the other eight cases in the table are boundary allocations.

And it's always useful to remember that a consumer's MRS at a bundle is the "personal value" one of the goods has to him, *measured in terms of another good*, *i.e.*, it tells us how much of the other good the consumer would be willing to give up to get a marginal increase in the good in question. This is extremely useful in trying to find Pareto improvements, and in seeing when no Pareto improvements are possible.

Table 1

	x_1^1	x_2^1	x_1^2	x_2^2		Required Relation between MRS's	Cases
(1)	$+$	$+$	$+$	$+$	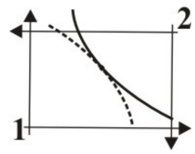	$MRS^1 = MRS^2$	(A) & (B)
(2)	0	$+$	$+$	$+$	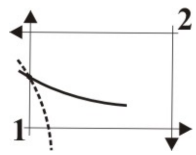	$MRS^1 \leqq MRS^2$	(B)
(3)	$+$	$+$	0	$+$	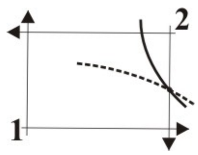	$MRS^1 \geqq MRS^2$	(A)
(4)	$+$	0	$+$	$+$	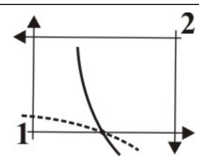	$MRS^1 \geqq MRS^2$	(A)
(5)	$+$	$+$	$+$	0	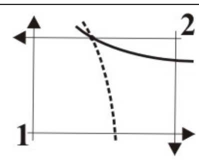	$MRS^1 \leqq MRS^2$	(B)
(6)	0	$+$	$+$	0	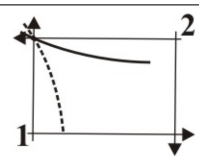	$MRS^1 \leqq MRS^2$	(B)
(7)	$+$	0	0	$+$	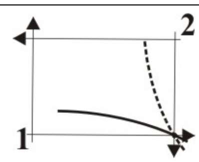	$MRS^1 \geqq MRS^2$	(A)
(8)	0	0	$+$	$+$	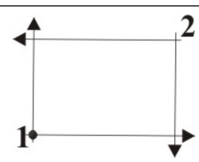	$-$	-
(9)	$+$	$+$	0	0	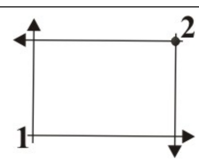	$-$	-

7 Maximizing a Social Welfare Function

An alternative approach for making welfare comparisons of alternative allocations is to evaluate the allocations according to a "social welfare function." We could in principle use any real-valued function W defined on the space $\mathbb{R}^{nl}_+$ of allocations $(\mathbf{x}^i)_1^n$. Of course, we would want to use a function that somehow reflects the preferences of the n consumers, so we'll define a social welfare function as any weighted sum of the consumers' utilities.

Definition: A **social welfare function** for the economy $(u^i, \mathring{\mathbf{x}}^i)_1^n$ is a function of the form $W(\mathbf{x}) = \sum_{i=1}^n \alpha_i u^i(\mathbf{x}^i)$ for some numbers (**weights**) $\alpha_1, ..., \alpha_n > 0$.

This may seem to be an ill-advised approach, because the social welfare function W adds up individual utilities that aren't really comparable: the consumers' utility functions have no cardinal meaning, because the underlying preferences can be represented by any monotone transforms of the given utility functions. But let's nevertheless see what the implications of using a social welfare function would be. Note that the map taking profiles of utility functions to a social welfare function, $(u^1, ..., u^n) \mapsto W(\cdot)$, is a particular way of aggregating profiles of utility functions into an "aggregate" utility function, as promised in Section 2. In keeping with our notation for aggregating preference relations, it would be natural to denote the social welfare function as $\bar{u}(\cdot)$; we use $W(\cdot)$ instead, because that's the conventional notation for a social welfare function.

The first thing we see is that any allocation that maximizes a social welfare function is Pareto efficient:

Theorem: If an allocation $\widehat{\mathbf{x}} \in \mathbb{R}^{nl}_+$ is a solution of the problem

$$\begin{aligned} \max_{(x_k^i) \in \mathbb{R}^{nl}_+} W(\mathbf{x}) = \textstyle\sum_{i=1}^n \alpha_i u^i(\mathbf{x}^i) \\ \text{subject to} \quad x_k^i &\geqq 0, \quad i = 1, ..., n, \quad k = 1, ..., l \\ \sum_{i=1}^n x_k^i &\leqq \mathring{x}_k, \quad k = 1, ..., l \end{aligned} \tag{W-Max}$$

for some numbers $\alpha_1, ..., \alpha_n > 0$, then $\widehat{\mathbf{x}}$ is a Pareto allocation for $\mathring{\mathbf{x}}$ and the utility functions $u^1, \ldots u^n$.

In fact, this result is much more general. It holds not just for our economic allocation problem, but for *any* situation in which we want to aggregate individual preferences into a single aggregate preference and in which the individual preferences can each be represented by a utility function. As the proof below makes clear, the result follows immediately from the definition of Pareto efficiency. As in the definition, the set X of alternatives here can be any set whatsoever.

Theorem: If the alternative $\widehat{x}$ is a solution of the problem

$$\max_{x \in X} W(x) = \sum_{i=1}^{n} \alpha_i u^i(x), \tag{5}$$

for some numbers $\alpha_1, ..., \alpha_n > 0$, then $\widehat{x}$ is Pareto efficient in X.

Proof. Suppose $\widehat{x}$ is *not* Pareto efficient in X: let $\widetilde{x}$ satisfy

$$\forall i \in N : u_i(\widetilde{x}) \geq u_i(\widehat{x}) \text{ and } \exists j \in N : u_j(\widetilde{x}) > u_j(\widehat{x}). \tag{6}$$

Then for any $\alpha_1, ..., \alpha_n > 0$ we have

$$\sum_{i \in N} \alpha_i u_i(\widetilde{x}) > \sum_{i \in N} \alpha_i u_i(\widehat{x}) \tag{7}$$

— *i.e.*, there are no values of the α_i for which $\widehat{x}$ maximizes $W(\cdot)$ on X, contrary to assumption. □

What about the converse? For any Pareto allocation $\widehat{\mathbf{x}}$, can we always find weights $\alpha_1, ..., \alpha_n$ for which $\widehat{\mathbf{x}}$ maximizes the social welfare function $\max_{(x_k^i) \in \mathbb{R}_+^{nl}} W(\mathbf{x}) = \sum_{i=1}^{n} \alpha_i u^i(\mathbf{x}^i)$? The answer is *no*; the following exercise asks you to construct a counterexample.

Exercise: In a two-person, two-good exchange economy, assume that $u_A(x_A, y_A) = x_A y_A$ and that $u_B(x_B, y_B) = x_B y_B$ and that the total resources are $\mathring{x}$ and $\mathring{y}$. Depict the set of Pareto allocations in the Edgeworth box. Then show that if $\alpha = \beta$ there are exactly two allocations that maximize the social welfare function $W(x_A, y_A, x_B, y_B) = \alpha u_A(x_A, y_A) + \beta u_B(x_B, y_B)$. Use this result, along with the corresponding result for $\alpha \neq \beta$, to establish that this example is indeed a counterexample to the converse of the above theorem.

Suggestion: Write r for the ratio $\mathring{y}/\mathring{x}$ and show that Pareto efficiency and maximization of W each require that $y_i = r x_i$ for $i = A, B$. This allows you to express u_A, u_B, and W in terms of just x_A and x_B, and now you can draw the constraint and the contours of W in the two-dimensional $x_A x_B$-space and easily establish the conclusion both geometrically and algebraically. Do it first for the case $\alpha = \beta$, where there are two (and only two) allocations that maximize W.