NO PRICES NO GAMES!

FOUR ECONOMIC MODELS

michael RICHTER
ariel RUBINSTEIN

© 2024 Michael Richter and Ariel Rubinstein

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs license (CC BY-NC-ND 4.0). This license allows you to share, copy, distribute, and transmit the work providing you do not modify the work, you do not use the work for commercial purposes, you attribute the work to the authors, and you provide a link to the license. Attribution should not in any way suggest that the authors endorse you or your use of the work and should include the following information.

Michael Richter and Ariel Rubinstein, *No Prices No Games!*. Second Edition. Cambridge, UK: Open Book Publishers, 2024, https://doi.org/10.11647/OBP.0438

Further details about CC BY-NC-ND licenses are available at https://creativecommons.org/licenses/by-nc-nd/4.0/.

All external links were active at the time of publication unless otherwise stated and have been archived via the Internet Archive Wayback Machine at https://archive.org/web.

ISBN Paperback: 978-1-80511-308-9
ISBN Hardback: 978-1-80511-309-6
ISBN Digital (PDF): 978-1-80511-310-2
DOI: 10.11647/OBP.0404

Information about any revised version of this work will be provided at https://doi.org/10.11647/OBP.0438.

Cover image: Ariel Rubinstein
Cover concept: Michael Richter and Ariel Rubinstein
Cover design: Jeevanjot Kaur Nagpal

Version: 2024.11.12

Check for updates

4 Biased Preferences Equilibrium

In any economy, the core tension is between agents' wants and societal feasibility, and an equilibrium notion finds a balance between them. In Chapters 2 and 3, we investigated equilibrium notions that invoke various social mechanisms to achieve that balance: norms emerge that affect agents' opportunity sets such that if every agent optimizes his preference relation, then the profile of optimal choices is feasible.

This chapter takes a different approach. We follow Rubinstein and Wolinksy (2022) who propose a solution concept which captures a different social mechanism that can resolve the fundamental conflict between wants and feasibility: *agents' preference relations are systematically biased*. The bias does not affect the agents' opportunity sets but, rather, their preferences, which are systematically biased in such a way that the profile of agents' *biased optimal choices* is feasible.

Recall Aesop's classic fable (translation from Gibbs (2002)):

> Driven by hunger, a fox tried to reach some grapes hanging high on the vine but was unable to, although he leaped with all his strength. As he went away, the fox remarked "Oh, you aren't even ripe yet! I don't need any sour grapes."

In this fable, there is one agent, the Fox, and two alternatives, "picking the grapes" and "not picking the grapes". The economic problem is that the Fox initially prefers the former alternative but only the latter alternative is feasible. The conflict in the fable is resolved not by restricting the Fox's opportunities but, rather, by biasing his preferences so that he now prefers not to pick the grapes (which in his mind are turned to "sour grapes").

©2024, Michael Richter and Ariel Rubinstein, CC BY-NC-ND 4.0
https://doi.org/10.11647/OBP.0438.04

Preference biases are not just a matter for fables. Introspection tells us that feasibility often influences our preferences in everyday life. We often assign greater value to what we can obtain (such as being an economist) and less to what we cannot (such as being a politician. Of course, there are also circumstances where the opposite is true, and the more unobtainable something is the more desirable it becomes.

The Fox biased his preferences, and harmony was achieved. Similarly, we envision biases as a mechanism for bringing harmony to a multi-agent economy. These biases, like prices, will be systematic and apply uniformly to all agents. Every agent's final preferences are determined by both the commonly shared bias and his initial preferences. Thus, in contrast to a competitive equilibrium where prices affect choice sets and preferences are fixed, in a biased preferences equilibrium biases affect preferences and choice sets are fixed. This illustrates the dual roles played by prices and preferences in standard economic settings.

This chapter models social situations in which preferences invisibly respond to feasibility pressures, just as price adjustments achieve harmony in a competitive market. This approach echoes one proposed by Gintis (1972) who criticizes the neoclassical general equilibrium approach because it takes preferences as exogenous, and who instead proposes modelling agents' preferences as endogenous and determined by their role in society.

Other papers treat preferences as being influenced by equilibrium actions. In some of those, the change in preferences is a side effect of an agent's own action (for example, smoking may influence the desire to smoke in the future, as modelled by Becker and Murphy (1988)). In others, the change in preferences is the outcome of a deliberate action by an interested party (for example, advertisers seek to influence customers' preferences to their own advantage, as modelled by Bagwell (2007)).

4.1 The Economy and the Equilibrium Concept

In this chapter, the notion of an economy is modified to accommodate modelling systematic preference biases.

Definition: An Economy

An **economy** is a tuple $\langle N, (X^i)_{i\in N}, K, ((u^i_k)_{k\in K})_{i\in N}, F\rangle$ where:

- For each agent i, X^i is his *fixed* personal choice set.
- The set K is a set of **considerations** common to all agents.
- For each agent i, $(u^i_k)_{k\in K}$ is a tuple of **consideration functions** over X^i such that i's utility function over X^i is $\Sigma_k u^i_k(x)$.
- The set of feasible profiles, F, is a subset of $\Pi_{i\in N} X^i$.

This definition modifies our notion of an economy in two ways. First, and less importantly, different agents can have different choice sets. This allows for modelling a variety of settings. For example, an exchange economy with a set of goods K, a fixed price vector p, and initial endowment profile (e^i) can be modelled by setting $X^i = \{x \in \mathbb{R}^K_+ \mid p\cdot x = p\cdot e^i\}$ and $F = \{(x^i) \mid \Sigma x^i = \Sigma e^i\}$. Another example is a two-sided matching market with two equally-sized populations A and B. This can be modelled by setting $X^i = B$ for any $i \in A$ and $X^j = A$ for any $j \in B$, while F is the set of all profiles (x^i) for which for every i, j, $x^i = j$ implies $x^j = i$.

The more important modification of the original definition of an economy is the use of a different notion of preferences. Rather than specifying an ordinal preference relation over the set of alternatives, we use a type of utility function that enables us to model systematic biases. All agents share the same set of considerations K. Each agent i is characterized not by an ordinal preference relation, but by a vector of consideration functions $u^i = (u^i_k)_{k\in K}$ where $u^i_k(x)$ represents the impact of consideration k on his overall evaluation of the alternative x. The consideration functions are not constant and, where applicable, are differentiable. Agent i's overall utility from an alternative x is the sum of the utilities obtained from those considerations, i.e. $\Sigma_{k\in K} u^i_k(x)$.

A preference bias is modelled as a systematic and uniform change in the weights placed on the considerations. Let $\mathfrak{B} = \mathbb{R}^K_{++}$ be the set of *biases*. A bias $\beta = (\beta_k)_{k\in K}$ transforms *every* vector of consideration functions $u = (u_k)_{k\in K}$ into the vector of biased consideration functions $T(u,\beta) = (\beta_k u_k)_{k\in K}$.

That is, if an agent i enters the model with the vector of consideration functions (u^i_k) and the bias vector is (β_k), then the agent behaves as if his vector of consideration functions is $(\beta_k u^i_k)$ and chooses from X^i by maximizing $\Sigma_{k\in K}\beta_k u^i_k(x)$. The preferences are unbiased when $\beta = (c,\ldots,c)$ because (u^i_k) and (cu^i_k) induce the same preferences. Any bias can be normalized to sum to 1 and therefore, can be naturally interpreted as a vector of weights on the different considerations.

The specification of K consideration functions is important beyond inducing an agent's preference relation. Two vectors of consideration functions, $u = (u_k)$ and $v = (v_k)$, can induce the *same* preference relation, and yet their biased preferences, $T(u,\beta)$ and $T(v,\beta)$, may induce *different* preference relations, as shown in the following example:

Let $K = \{1,2\}$, $X^1 = X^2 = \{0,1\}\times\{0,1\}$ (the four corners of the unit square), $u_1(x) = x_1$, $u_2(x) = 4x_2$, and $v_1(x) = x_1$, $v_2(x) = 2x_2$. Both (u_k) and (v_k) induce the same preferences $(1,1) \succ (0,1) \succ (1,0) \succ (0,0)$, but when the bias vector is $\beta = (3,1)$, the preferences induced by $T((u_k),\beta)$ are unchanged while the preferences induced by $T((v_k),\beta)$ become $(1,1) \succ (1,0) \succ (0,1) \succ (0,0)$.

On the other hand, when the set of alternatives X is the positive orthant of an Euclidean space, any two additively separable, monotonic, and differentiable utility functions that represent the same preferences on X are transformed by the bias map T into the same biased preferences.

Claim: Preference Preservation in Euclidean Spaces

Let $X = \mathbb{R}^{|K|}_+$ be the common set of alternatives. Let $u = (u_k)_{k\in K}$ and $v = (v_k)_{k\in K}$ be vectors of consideration functions, so that both u_k and v_k depend only on x_k and are differentiable with strictly positive derivatives. If u and v induce identical preference relations, then for any β, the biased preferences induced by $T(u,\beta)$ and $T(v,\beta)$ are identical.

Proof:

The case $K = \{1\}$ is vacuous since the only admissible preference relation is the increasing ordering on $\mathbb{R}_+$.

Let $|K| \geq 2$. Since u and v represent the same preferences, then at any $x \in X$ the gradient of the utility function $V(x) = \Sigma v_k(x_k)$ is a rescaling of the gradient of $U(x) = \Sigma u_k(x_k)$. That is, for every x, there is a strictly positive scalar $\mu(x) > 0$ such that $\nabla V(x) = \mu(x)\nabla U(x)$. It suffices to show that $\mu(x)$ is a constant μ since then $v'_k(x_k) = \mu u'_k(x_k)$ for all k and x_k, and therefore $v_k(x_k) = \mu u_k(x_k) + c_k$ for some c_k. Thus, $\Sigma\beta_k v_k(x_k) = \Sigma\beta_k(\mu u_k(x_k) + c_k) = \mu\Sigma\beta_k u_k(x_k) + \Sigma\beta_k c_k$ and therefore, the biased utility functions are an affine transformation of each other and, as such, represent the same preferences over X.

To show that $\mu(x)$ is a constant, notice that if two bundles x and y share a coordinate $x_k = y_k$, then $\mu(x) = \mu(y)$ since $v'_k(x_k) = \mu(x)u'_k(x_k)$ and $v'_k(y_k) = \mu(y)u'_k(y_k)$. If they do not share a coordinate, then let z be a bundle such that $z_1 = x_1$ and $z_2 = y_2$. Since z shares a coordinate with each of them, it must be that $\mu(x) = \mu(z) = \mu(y)$.

Generalizing the bias concept. The bias notion has been formalized for preference relations induced from an additively separable representation $\Sigma_k u_k(x)$. In particular, when X is a subset of a Euclidean space and u_k is a function only of x_k, a bias vector (β_k) multiplies the subjective tradeoff between any two goods k and l by β_k/β_l.

This suggests a generalization of the bias concept to more general preference relations where the biased preference tradeoffs are obtained by systematically multiplying the agent's unbiased tradeoffs. This generalization can be shown when $K = 2$ and preferences are differentiable. More precisely, for any preference relation $\succsim$, denote the marginal rate of substitution between k and l at a bundle x by $MRS_{k,l}(x, \succsim)$. For every vector of biases (β_1, β_2), there is another preference relation $\succsim'$ for which $MRS_{1,2}(x, \succsim') = \frac{\beta_1}{\beta_2} MRS_{1,2}(x, \succsim)$ for every bundle x. However, this is not always possible when $K > 2$:

Example: The Difficulty in Extending the Bias Notion

Let $K = \{1,2,3\}$ and $\succsim$ be represented by u where $u(x_1, x_2, 0) = 2x_1$ and $u(x_1, x_2, 1) = x_1 + x_2$. Let the bias vector be $\beta = (\beta_k) = (1,2,1)$. Suppose that there are biased preferences $\succsim^\beta$ such that for any two goods k and l, $MRS_{k,l}(x, \succsim^\beta) = \frac{\beta_k}{\beta_l} MRS_{k,l}(x, \succsim)$ at every alternative x. Then:

(i) The $MRS_{1,2}$ is unchanged (and remains ∞) at every $(x_1, x_2, 0)$.

(ii) The $MRS_{1,2}$ is changed from 1 to $1/2$ at every $(x_1, x_2, 1)$.

(iii) The $MRS_{1,3}$ is unchanged at any alternative.

By (i) and (iii), it holds that $(3,5,1) \sim^\beta (4,5,0) \sim^\beta (4,6,0) \sim^\beta (2,6,1)$ because $u(3,5,1) = u(4,5,0) = u(4,6,0) = u(2,6,1)$. But then, $(3,5,1) \sim^\beta (2,6,1)$, contradicting (ii). Thus, there is no preference relation $\succsim^\beta$ for which all biased marginal rates of substitution are β-scaled versions of the original.

We now define the solution concept. As explained above, harmony will be achieved by a uniform social shift of the weights placed on the consideration functions. An equilibrium consists of a profile (x^i) of alternatives and a bias β such that: (i) for every agent i, the alternative x^i is optimal in X^i according to his biased preferences; and (ii) the profile is feasible.

Definition: Biased Preferences Equilibrium

A *biased preferences equilibrium* is a tuple $\langle \beta, (x^i) \rangle$ where $\beta \in \mathfrak{B}$ and (x^i) is a profile of choices, such that:

(i) For every agent i, the alternative x^i is optimal in X^i according to the preferences induced by $T(u^i, \beta)$.

(ii) The profile (x^i) is in F.

We will refer to a profile of alternatives that is Pareto-optimal according to the agents' initial preferences as *pre-Pareto optimal*. Obviously, since agents choose optimally given their biased preferences, a profile of biased preferences equilibrium choices is always ex-post Pareto optimal.

Example: The Service Economy

Let $X = \{a, b\}$ where $x \in X$ is a charging station. The amount of time that agent i needs to charge his vehicle is $t^i \leq 1$. The charging stations vary along two aspects a and b with each aspect possessed by only one of the two stations. That is, the utility function of agent i satisfies $v_a^i(a) > 0$, $v_a^i(b) = 0$ and $v_b^i(a) = 0$, $v_b^i(b) > 0$. Each station can serve agents as long as the total time required to serve them is not more than 1, thus F includes all the profiles (x^i) such that $\sum_{\{i|x^i=x\}} t^i \leq 1$ for both $x \in X$.

Let $\alpha^i = v_a^i(a)/v_b^i(b)$ and assume that $\alpha^1 > \alpha^2 > \cdots > \alpha^n$. Given a bias vector β, denote $\mu = \beta_b/\beta_a$. An agent i will choose a if $\alpha^i > \mu$, he will choose b if $\alpha^i < \mu$, and will be indifferent between the options if $\alpha^i = \mu$.

The following service economy has a unique biased preferences equilibrium profile. Agents $1, \ldots, 4$ initially prefer to be served in a while agent 5 initially prefers to be served at b.

Agent	1	2	3	4	5
α^i	5	4	3	2	1/3
t^i	0.3	0.6	0.7	0.1	0.1

Table 4.1 An equilibrium in the Service economy

The biased preferences equilibrium profile is unique: agents 1 and 2 choose a while the others choose b and is supported by any bias for which $3 \leq \mu \leq 4$. This is not a pre-Pareto optimal profile: moving agent 4 from b to a is a Pareto improvement that is impossible in the biased preferences equilibrium because blocking agent 3 from attending a forces the bias to be such that agent 4 is also biased towards b.

Of course, an equilibrium may not exist. For example, if the above table is modified so that $t^5 = 0.25$ (instead of 0.1), then no biased preferences equilibrium exists.

In the rest of the chapter, we analyze the biased preferences equilibrium for several examples: the give-and-take economy, the exchange economy with fixed prices, and a few housing-type economies.

4.2 The Give-and-Take Economy

We return to an old friend: the give-and-take economy. Recall that, in this economy, each agent decides how much to contribute to or withdraw from a social fund, and feasibility requires that the total contributions equal the total withdrawals. All agents face the same choice set $[-1,1]$, where, as usual, a positive number represents the amount taken from (and a negative number the amount given to) the social fund. To fit it into the current framework, let the two considerations be g (giving) and t (taking). The consideration g is generous (people like to give), while the consideration t is selfish (people like to take). Initially, an agent i's choice balances between these two considerations by choosing the maximizer of $u_g^i(x)+u_t^i(x)$, where u_g^i is a strictly decreasing and strictly concave function while u_t^i is a strictly increasing and strictly concave function. It follows that both are continuous. Denote by $peak^i$ the unique maximizer of $u_g^i(x)+u_t^i(x)$, which is agent i's most-preferred choice in $[-1,1]$, and assume that it is interior.

Any bias β pushes all agents' preferences in the same direction by systematically altering the tradeoff between generosity and selfishness. Denote $\mu=\beta_t/\beta_g$. A μ above 1 biases the agents' preferences towards selfishness, while a μ below 1 biases the agents' preferences towards generosity.

Under the above assumptions, there is a unique equilibrium profile, and it is pre-Pareto optimal:

Proposition 4.1: Uniqueness and Pre-Pareto Optimality

(i) The give-and-take economy has a unique biased preferences equilibrium (up to a rescaling of the bias vector).

(ii) The equilibrium profile is pre-Pareto optimal.

Proof:

(i) Given the bias $(1,\mu)$, each agent i optimizes $u_g^i(x)+\mu\cdot u_t^i(x)$ and has a unique optimal choice denoted by $x^i(\mu)$. Since all of the x^i functions are increasing and continuous, so is the net overall "demand" from the social fund, $\Sigma_i x^i(\mu)$. This sum is positive when μ is sufficiently large and negative when μ is sufficiently small. Therefore, there is a μ^* for which the sum is zero, and $\langle(1,\mu^*),(x^i(\mu^*))\rangle$ is a biased preferences equilibrium.

The parameter μ^* is unique since $x^i(\mu)$ is strictly increasing when $-1<x^i(\mu)<1$, and it must be that $x^i(\mu^*)$ is interior for some i. (It is impossible that all agents choose 1 or all choose -1. It is also impossible that in equilibrium some choose 1 and the other -1 since if $\mu^*\geq 1$, then for every i, $x^i(\mu^*)\geq \mathit{peak}^i > -1$ and if $\mu^*\leq 1$, then for every i, $x^i(\mu^*)\leq \mathit{peak}^i < 1$.)

(ii) If $\mu^*=1$, then every agent's equilibrium choice is his unbiased first-best. If $\mu^*>1$, then every agent i chooses $x^i(\mu^*)\geq \mathit{peak}^i$, and any ex-ante Pareto improvement (y^i) must satisfy $y^i\leq x^i(\mu^*)$ with at least one strict inequality, but that contradicts feasibility since $0=\Sigma x^i(\mu^*)>\Sigma y^i$. Thus, the equilibrium profile is pre-Pareto optimal. The case $\mu^*<1$ is analogous.

Note that in the biased preferences equilibrium, balancing the social fund is a shared responsibility of all agents: when the sum of the agents' peaks is positive (i.e. there is an overall preference for taking), the equilibrium bias $(\mu^*<1)$ overweighs generosity and *every* agent chooses a point below his peak. This is essentially true for primitive equilibria (in every equilibrium profile all agents are assigned to alternatives weakly below their peaks). In contrast, in a Y-equilibrium, there is a uniform cap on withdrawals and only the greediest agents are impacted, and in a jungle equilibrium, only the weakest agents are restricted.

4.3 The Fixed-Prices Exchange Economy

The next example is related to the literature on economies with fixed prices (see Benassy (1986) and the references therein). Let $X = \mathbb{R}^K_{++}$ be the set of bundles in a world with a set of goods K. Every agent i has an initial endowment $e^i \in X$, and exchange takes place according to a fixed price vector $p = (p_k)$. Accordingly, $X^i = \{x \in X \mid p \cdot x = p \cdot e^i\}$, and the set of feasible profiles is $F = \{(x^i) \in \Pi_i X^i \mid \Sigma_i x^i = \Sigma_i e^i\}$. All agents share the same considerations, one for each good. Each consideration function $u^i_k(x)$ is a function of only x_k, which is assumed to be increasing, twice-differentiable, and strictly concave.

In economies with fixed prices, rationing is typically the mechanism used to achieve harmony. That is, upper bounds are established on the consumable quantity of each good. In contrast, in a biased preferences equilibrium, economic harmony is achieved by means of a systematic adjustment of preferences.

Proposition 4.2: Biased Preferences Equilibria in Exchange Economies with Fixed Prices

In any exchange economy with fixed prices:

(i) A biased preferences equilibrium exists.

(ii) All biased preferences equilibrium outcomes are pre-Pareto optimal.

Proof:

(i) To illustrate, consider the two-good two-agent case, which can be depicted using an Edgeworth Box (see Figure 4.1). Assume that agents do not like consuming on the boundary, i.e. the derivative of every consideration function u^i_k at 0 is infinity. Let $(x^1, x^2) \in X^1 \times X^2$ be an overall Pareto-optimal allocation of $e^1 + e^2$, which always exists. Then, at (x^1, x^2), both agents have the same marginal rate of substitution μ. If $\mu = p_1/p_2$, then no bias is needed, that is, $\langle \beta = (1,1), (x^i) \rangle$ is a biased preferences equilibrium. If $\mu \neq p_1/p_2$, then the bias $\beta = (p_1, p_2\mu)$

modifies both preferences so that the $MRS_{1,2}$ of the biased preferences of each i at x^i is $\mu\beta_1/\beta_2 = p_1/p_2$. Thus, $\langle\beta,(x^i)\rangle$ is a biased preferences equilibrium.

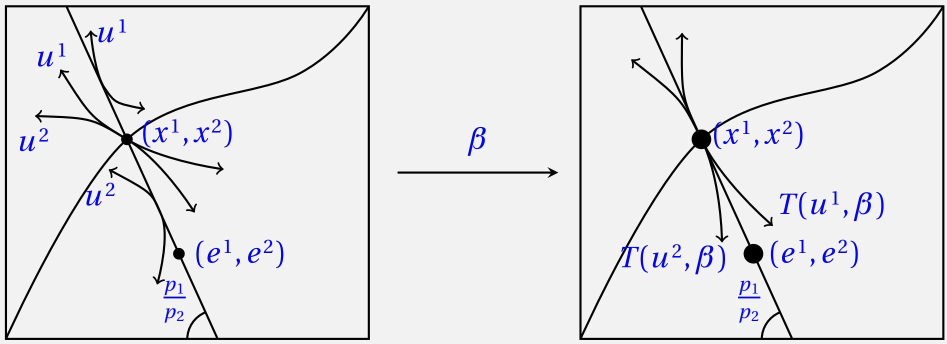

Figure 4.1 Equilibrium in an Edgeworth Box

For the case of any number of agents and more than two goods (and even without the boundary assumptions), Keiding (1981) (following Balasko (1979)) showed that there is a vector $q=(q_k)$ and an allocation (x^i) such that $p\cdot x^i = p\cdot e^i$ for all i, and if $y^i \succ^i x^i$, then $q\cdot y^i > q\cdot x^i$. Therefore, for every agent i, any good k that he consumes and any other good l, it holds that $MRS_{k,l}(x^i)\geq q_k/q_l$. Consequently, by setting $\beta=(p_k/q_k)_{k\in K}$, it holds that $biasedMRS_{k,l}(x^i) = \frac{\beta_k}{\beta_l}MRS_{k,l}(x^i) \geq \frac{\beta_k q_k}{\beta_l q_l} = \frac{p_k}{p_l}$ (the bound is an equality if $x^i_l > 0$). Therefore, for every agent i, given the price vector p and the initial bundle e^i, the bundle x^i is optimal for i's biased preferences. Thus, $\langle\beta,(x^i)\rangle$ is a biased preferences equilibrium.

(ii) Let $\langle\beta,(x^i)\rangle$ be a biased preferences equilibrium. Then, for each agent i and any good l that he consumes, the $MRS_{k,l}$ of the biased preferences $T(\beta,u^i)$ at x^i is bounded from above by p_k/p_l. Therefore, the $MRS_{k,l}$ of his initial preferences at x^i is bounded from above by $\frac{p_k/\beta_k}{p_l/\beta_l}$. Thus, (x^i) is a Walrasian equilibrium outcome in the unbiased economy with price vector $\left(p_k/\beta_k\right)$ and initial endowment (x^i). Therefore, by the standard First Welfare Theorem, (x^i) is pre-Pareto optimal.

We now consider an example with two goods and linear preferences where the biased preferences equilibrium can easily be calculated.

Example: Linear Preferences

Suppose that there are two agents, two goods, and that for every agent i, the two consideration functions are linear, that is, $u_1^i(x_1) = x_1$ and $u_2^i(x_2) = \alpha^i x_2$ where every α^i is a positive number. Consider the configuration depicted in Figure 4.2:

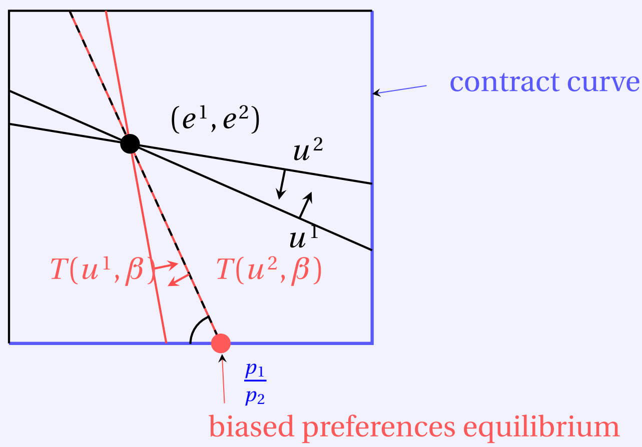

Figure 4.2 Biased linear preferences in an Edgeworth Box (dashed line = the budget lines; black solid lines = initial preferences; red solid lines = biased preferences; blue line = the contract curve for the initial preferences)

In this example (other configurations can be analyzed similarly):

(i) Agent 2 likes good 2 more than agent 1 does, that is, $\alpha^2 > \alpha^1$.

(ii) The ratio p_1/p_2 is greater than both agents' (constant) personal marginal rates of substitution, that is, $p_1/p_2 > 1/\alpha^1 > 1/\alpha^2$.

(iii) In any feasible allocation, both agents must consume positive amounts of good 1.

The economy is not in harmony because, given (ii), both agents wish to purchase only good 2.

In any biased preferences equilibrium, there is an agent i who consumes good 2, and by (iii), he also consumes good 1. By the linearity of the preferences, agent i must be indifferent between all alternatives in X^i, that is, $p_1/p_2 = \beta_1/(\alpha^i \beta_2)$. If $i = 1$, then by (i), agent 2 does not consume good 1, violating (iii). Thus, it must be that $i = 2$, and the bias satisfies $p_1/p_2 = \beta_1/(\alpha^2 \beta_2)$.

Such a bias is part of the equilibrium depicted in Figure 4.2, where agent 1 consumes only good 1, and agent 2 (who is indifferent between all bundles in his budget set) consumes all of good 2 and the remainder of good 1. It follows that this is the unique biased preferences equilibrium.

Failure of Individual Rationality: An interesting feature of a biased preferences equilibrium is that, even though it is pre-Pareto optimal, "Individual Rationality" can fail: in an equilibrium, an agent might choose a bundle that is inferior to his endowment bundle when judged by his initial preferences, as in the previous example. By his original preferences, agent 1 is worse off in the equilibrium than he was with his initial endowment, since he trades some of his good 2 endowment for good 1, but ex-ante he would prefer to do the opposite.

Example: Non-Convex Preferences

In the standard exchange economy with non-convex preferences, a competitive equilibrium may not exist: there may be no price vector for which the sum of the demands equals the total bundle. Nevertheless, there may be a price vector for which a biased preferences equilibrium exists. Thus, prices and biased preferences together may achieve harmony when the standard competitive equilibrium tools fail to do so.

To illustrate, consider the division economy where both agents have the non-convex preferences represented by $(x_1)^2 + 2(x_2)^2$ and the initial endowments are $e^1 = (1,1)$ and $e^2 = (2,2)$. There is no standard competitive equilibrium. Given any price vector, each agent will consume only one of the two goods, and since the agents have the same preferences, any equilibrium price vector must make each agent indifferent between the two goods, i.e. $p = (1, \sqrt{2})$. But, then agent 2 will demand more than 3 units of one of the goods.

In contrast, a biased preferences equilibrium exists. Let $p = (2, 1)$ and $\beta = (8, 1)$. Each agent's biased utility function is $4(x_1)^2 + (x_2)^2$. Agent 1's optimal bundles are $(1.5, 0)$ and $(0, 3)$, and agent 2's optimal bundles are $(3, 0)$ and $(0, 6)$. Thus, the bias β, together with the allocation $x^1 = (0, 3)$ and $x^2 = (3, 0)$, is a biased preferences equilibrium in the exchange economy with fixed prices p.

Note that agent 1 is initially poorer than agent 2, but in the equilibrium, agent 1 is better off according to the initial preferences!

4.4 Housing-Type Economies

We return to the classic housing economy of Shapley and Scarf (1974), in which there is a set N of agents and an equally-sized set H of houses. Each agent i chooses a single house, that is, $X^i = H$. Let $v^i(h) > 0$ be agent i's valuation of house h. The model can be enriched to fit our framework by taking the set of considerations to be H and setting $u^i_h(x^i) = v^i(h)$ if $x^i = h$ and 0 otherwise. Given a bias vector (β_h), an agent i derives utility $\beta_h v^i(h)$ from house h.

Example:

The following table presents the consideration function values in a housing economy with two agents.

	h_1	h_2
$v^1(h)$	4	3
$v^2(h)$	3	1

Table 4.2 House utilities

Both agents initially prefer house h_1. To achieve harmony, the bias must boost h_2 so that one agent will choose it, but not to the extent that both will. For example, a biased preferences equilibrium is obtained by the bias $(1,2)$, which results in agent 1 choosing h_2 and agent 2 choosing h_1. Of course, other biases are possible but, in all biased preferences equilibria, agent 1 gets h_2 and agent 2 gets h_1. Note that, in the biased preferences equilibrium profile, the product of the ex-ante values $(3 \cdot 3 = 9)$ is larger than that in the other assignment $(4 \cdot 1 = 4)$. We will see below that this is not a coincidence.

We say that a feasible profile (x^i) is *Nash maximal* if it maximizes $\Pi_{i \in N} v^i(x^i)$ over all feasible profiles. We now show that the set of biased preferences equilibrium profiles is precisely the set of Nash-maximal profiles and thus, any biased preferences equilibrium profile is pre-Pareto optimal. The proof is a direct application of Shapley and Shubik (1971) (see also Gale (1984) for a proof using the KKM Lemma).

Proposition 4.3: Biased Preferences Equilibrium = Nash Maximality

In the housing economy, the set of biased preferences equilibrium profiles is the set of Nash-maximal profiles.

Proof:

Let $(h^i)_{i \in N}$ be a Nash-maximal profile, that is, it maximizes $\Sigma_{i \in N} \ln(v^i(x^i))$ over all feasible assignments. By Shapley and Shubik (1971), there exists a price vector (p_h) so that for each agent i, the house h^i is a maximizer of $\ln(v^i(x^i)) - p_{x^i}$, and therefore, it is also a maximizer of $v^i(x^i)/e^{p_{x^i}}$. Thus, $\langle (\beta_h = 1/e^{p_h})_{h \in H}, (h^i)_{i \in N} \rangle$ constitutes a biased preferences equilibrium.

In the other direction, let $\langle \beta, (h^i) \rangle$ be a biased preferences equilibrium and (x^i) be any other assignment. For each i, $\beta_{h^i} v^i(h^i) \geq \beta_{x^i} v^i(x^i)$ and therefore, $\Pi_i \beta_{h^i} v^i(h^i) \geq \Pi_i \beta_{x^i} v^i(x^i)$. Since $\Pi_i \beta_{h^i} = \Pi_i \beta_{x^i}$, it follows that $\Pi_i v^i(h^i) \geq \Pi_i v^i(x^i)$, that is, (h^i) is Nash maximal.

We proceed with two modifications to the housing economy:

Example: The Partnership Economy

As in Shapley and Shubik (1971), the agents are composed of two equally-sized populations, A and B. Each agent chooses a unique partner from the other population, that is, $X^i = B$ for any $i \in A$ and $X^j = A$ for any $j \in B$. A profile is feasible if for every i, j, if i chooses j, then j chooses i. An agent i's valuation of a partnership with j is $v^i(j) > 0$. Importantly, the ex-ante valuations are assumed to be symmetric, that is, $v^i(j) = v^j(i)$, but the biased valuations might not be.

Example: Let $A = \{1,2\}$ and $B = \{3,4\}$. Table 4.3 presents the original valuations (left bi-matrix) and the equilibrium biased valuations (right bi-matrix). Each cell gives the values of i and j of being matched. The Nash-maximal matching is 1↔4 and 2↔3 (depicted), which is a biased preferences equilibrium with the bias $\beta = (2,1,2,1)$.

	3	4
1	1,1	3,3
2	3,3	4,4

$\xrightarrow{T(\cdot,\beta)}$

	3	4
1	2,2	3,6
2	6,3	4,4

Table 4.3 A Biased Preferences Equilibrium

Claim: In the partnership economy, the set of biased preferences equilibrium outcomes is the set of all Nash-maximal profiles.

Proof: Let (a^i) be a Nash-maximal profile, that is, one that maximizes $\Sigma_{i\in N} \ln(v^i(x^i))$ over F. Since the utility functions are symmetric, i.e. $(v^i(j) = v^j(i))$, the profile $(a^i)_{i\in A}$ maximizes $\Sigma_{i\in A} \ln(v^i(x^i))$ over the assignments of the members of A to B, and the profile $(a^i)_{i\in B}$ maximizes $\Sigma_{i\in B} \ln(v^i(x^i))$ over the assignments of the members of B to A. Taking the agents to be A and the houses to be B, Shapley and Shubik (1971) showed that a price vector $(p_j)_{j\in B}$ exists, such that for every agent $i \in A$, the choice of a^i maximizes $\ln(v^i(j)) - p_j$ over all $j \in X^i = B$ and therefore, maximizes $v^i(j)/e^{p_j}$ as well. Reversing roles, there is a price vector $(p_j)_{j\in A}$ with analogous optimality properties. Therefore, $\langle(\beta_j = 1/e^{p_j})_{j\in N}, (a^i)_{i\in N}\rangle$ constitutes a biased preferences equilibrium.

In the other direction, let $\langle(\beta_j)_{j\in N}, (a^i)_{i\in N}\rangle$ be a biased preferences equilibrium, and let $(x^i) \in F$. For every i, it holds that $\beta_{a^i} v^i(a^i) \geq \beta_{x^i} v^i(x^i)$. Therefore, $\Pi_i[\beta_{a^i} v^i(a^i)] \geq \Pi_i[\beta_{x^i} v^i(x^i)]$. Since $\Pi_i \beta_{a^i} = \Pi_i \beta_{x^i}$, it follows that $\Pi_i v^i(a^i) \geq \Pi_i v^i(x^i)$. That is, (a^i) is Nash maximal. ■

The condition that the value of a match between any two agents is the same for both of them is sufficient for the A-Nash-maximal matching to be B-Nash-maximal as well. Without this condition, a biased preferences equilibrium may not exist:

Example: Consider an assignment economy where $A = \{1,2\}$, $B = \{3,4\}$, and $3 \succ^1 4$, $1 \succ^4 2$, $4 \succ^2 3$, $2 \succ^3 1$. No utility presentation of these preferences is consistent with the assumption that the value of a match is identical for both partners (since it requires that $v^1(3) > v^1(4) = v^4(1) > v^4(2) = v^2(4) > v^2(3) = v^3(2) > v^3(1) = v^1(3)$).

Suppose that $1 \leftrightarrow 3$ is a match in a biased preferences equilibrium. Then, $\beta_1 > \beta_2$ (so that agent 3 chooses agent 1 over agent 2). But then, agent 4 will also choose agent 1, violating feasibility. Likewise, $1 \leftrightarrow 4$ cannot be a match in a biased preferences equilibrium.

Example: A Production Economy

In the production economy (related to Atakan, Richter, and Tsur (2023)) there is a set of indivisible goods K and two equally-sized groups of agents: *consumers* (C) and *producers* (P). Every $i \in C$ consumes exactly one unit of a single good, that is $X^i = K$, and $u^i_k > 0$ is consumer i's utility from consuming good k (which he wishes to *maximize*). Every producer $i \in P$ *must* produce exactly one unit of a single good, that is, $X^i = K$, and $c^i_k > 0$ is producer i's *utility-cost* from producing good k (which he wishes to *minimize*). The set F consists of all profiles satisfying that, for every good k, the number of its consumers is equal to the number of its producers. Note that this economy differs from the partnership economy in that consumers and producers choose a good rather than a partner and the biases are applied to the goods rather than to the agents.

A bias vector $\beta = (\beta_k)_{k\in K}$ alters consumer i's utility vector from $(u^i_k)_{k\in K}$ to $(\beta_k u^i_k)_{k\in K}$ and producer i's utility-cost vector from $(c^i_k)_{k\in K}$ to $(\beta_k c^i_k)_{k\in K}$. Thus, a bias β simultaneously rescales both the consumers' utility and the producers' utility-costs of good k by the same factor β_k. Thus, an *increase* in β_k is analogous to that of a *decrease* in the price of good k in a regular exchange economy: it makes the good more desirable to buyers and less desirable to sellers. Underlying a bias could be some trait such as quality: a high bias, like a high quality level, makes the good more desirable to consumers and increases the utility-cost to produce it.

The following claim again uses a Shapley and Shubik (1971)-style argument to characterize the biased preferences equilibrium profiles. It implies that they exist and are pre-Pareto optimal.

Claim: In the production economy, the biased preferences equilibrium profiles are precisely the solutions of:

$$\max_{(x^i)\in F} \frac{\prod_{i\in C} u^i_{x^i}}{\prod_{i\in P} c^i_{x^i}} \tag{*}$$

Proof: Let $\langle\beta,(x^i)\rangle$ be a biased preferences equilibrium, and let $(y^i)\in F$. It follows that $\beta_{x^i} u^i_{x^i} \geq \beta_{y^i} u^i_{y^i}$ for every $i \in C$ and $\beta_{x^i} c^i_{x^i} \leq \beta_{y^i} c^i_{y^i}$ for every $i \in P$. Combined with the equality $\Pi_{i\in C}\beta_{z^i} = \Pi_{i\in P}\beta_{z^i}$, which holds for all $(z^i)\in F$, we conclude that:

$$\frac{\prod_{i\in C} u^i_{x^i}}{\prod_{i\in P} c^i_{x^i}} = \frac{\prod_{i\in C} \beta_{x^i} u^i_{x^i}}{\prod_{i\in P} \beta_{x^i} c^i_{x^i}} \geq \frac{\prod_{i\in C} \beta_{y^i} u^i_{y^i}}{\prod_{i\in P} \beta_{y^i} c^i_{y^i}} = \frac{\prod_{i\in C} u^i_{y^i}}{\prod_{i\in P} c^i_{y^i}}$$

and therefore (x^i) is a solution of (*).

In the other direction, let (x^i) be a solution of (*). Let (x^i_k) be the allocation matrix with a row for each agent and a column for each good, where $x^i_k = 1$ if i chooses k and $x^i_k = 0$ otherwise. The matrix solves the following linear maximization problem:

$$\max_{(m^i_k)} \Sigma_k \Sigma_{i\in C}[\ln(u^i_k)m^i_k] + \Sigma_k \Sigma_{i\in P}[-\ln(c^i_k)m^i_k]$$

$$\text{such that} \qquad \Sigma_{i\in C} m^i_k - \Sigma_{i\in P} m^i_k = 0 \ \forall k \qquad (\mu_k)$$

$$m^i_k \geq 0 \ \forall i,k \qquad (\gamma^i_k)$$

$$\Sigma_k m^i_k = 1 \ \forall i \qquad (\psi^i)$$

The above problem always has a solution which is a binary matrix (that is, $x^i_k = 0$ or 1, for every i,k). To understand why, Birkhoff (1946) (and his extensions in Budish, Che, Kojima, and Milgrom (2013)) shows that any matrix of real numbers that satisfies the above constraints is a convex combination of binary matrices that also satisfy them. Since the target function is linear, any such binary matrix is also a solution to the linear programming problem.

The constraints in the above optimization are labelled by their shadow values, which appear in the parenthesis to the right. Let $\beta = (e^{\mu_k})_{k\in K}$. We will now verify that $\langle\beta,(x^i)\rangle$ is a biased preferences equilibrium.

The matrix (x^i_k) satisfies the following conditions for every $i \in C$ and $k \in K$:

$$\ln(u^i_k) + \mu_k + \gamma^i_k + \psi^i = 0,\ \gamma^i_k x^i_k = 0 \text{ and } \gamma^i_k \leq 0$$

If $x^i_k = 1$, then for any $k' \in K : \ln(u^i_k) + \mu_k + \psi^i = 0 \geq \ln(u^i_{k'}) + \mu_{k'} + \psi^i$ and therefore, $\beta_k u^i_k \geq \beta_{k'} u^i_{k'}$, that is, k is i's most preferred good according to the β-biased preferences. Likewise, for every $i \in P$ and $k \in K$:

$$-\ln(c^i_k) - \mu_k + \gamma^i_k + \psi^i = 0,\ \gamma^i_k x^i_k = 0 \text{ and } \gamma^i_k \leq 0$$

If $x^i_k = 1$, then for any k': $-\ln(c^i_k) - \mu_k + \psi^i = 0 \geq -\ln(c^i_{k'}) - \mu_{k'} + \psi^i$ and therefore, $c^i_k \beta_k \leq c^i_{k'} \beta_{k'}$, that is, k minimizes i's β-biased utility-costs. ■

Example: Table 4.4 (left panel) depicts a production economy with two goods, two consumers (C1 and C2), and two producers (P1 and P2). The right panel indicates a biased preferences equilibrium with $\beta = (1,3)$ where Consumer 1 and Producer 1 choose Good 1, and Consumer 2 and Producer 2 choose Good 2.

	Good 1	Good 2
C1	4	1
C2	4	2
P1	2	1
P2	4	1

$\xrightarrow{T(\cdot,\beta)}$

	1· Good 1	3· Good 2
C1	4	3
C2	4	6
P1	2	3
P2	4	3

Table 4.4 A Biased Preferences Equilibrium in a Production Economy.
The left panel presents the initial utilities and utility-costs for the agents and the right panel presents the biased utilities and utility-costs with the bias $\beta = (1,3)$.

Initially, both consumers prefer to consume good 1, while both producers prefer to produce good 2. In the solution of (*) (with the value $(4 \cdot 2)/(2 \cdot 1)$), C1 and P1 choose good 1, and C2 and P2 choose good 2. This outcome is obtained in equilibrium by, for example, $\beta = (1,3)$, a bias that induces one consumer to switch from good 1 to good 2, and one producer to switch from good 2 to good 1.