BEEM103 – Optimization Techiniques for Economists	Dieter Balkenborg Departments of Economics
Lecture Week 5	University of Exeter

1 Constrained Optimization

A constrained optimization problem (with two unknowns) consists of
- an *objective function*

$$z = f(x, y)$$

which is to be maximized (or to be minimized)
- a list of *constraints* which can be written in the form

$$\begin{aligned} g_1(x,y) &\geq 0 \\ g_2(x,y) &\geq 0 \\ &\vdots \\ g_K(x,y) &\geq 0 \end{aligned} \qquad (*)$$

whereby the objective function f and the *constraining functions* $g_1, \cdots, g_K$ are defined on the same domain D.

Example 1 The following graph illustrates the idea. The constraints $g_k(x,y) \leq 0$ carve out a region G where the optimum has to lie. A line indicates in which area the first constraint is satisfied ($g_1(x,y) \leq 0$) and where it is violated ($g_1(x,y) > 0$). We see some level curves of the objective function. The unconstrained optimum would be at the point A while constrained to the region G the optimum is at point B.

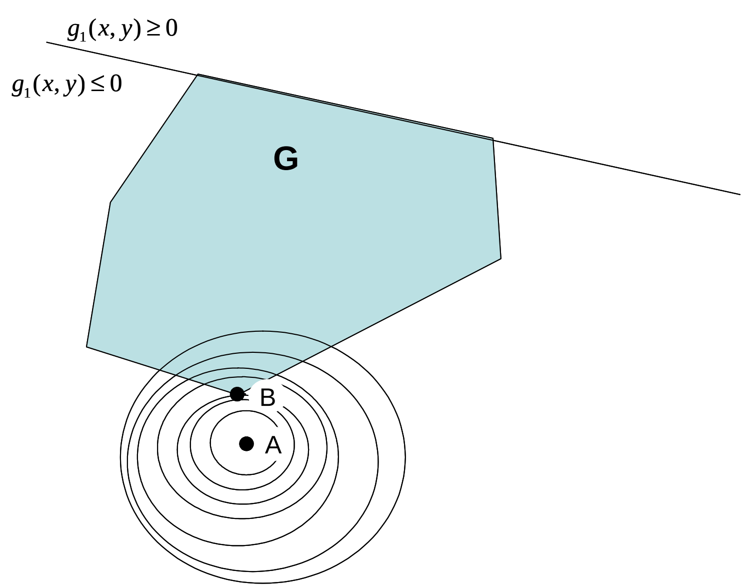

Remark 1 We consider here only inequality constraints. An equality constraint $g(x,y) = 0$ can be dealt with in this framework by splitting it up into two inequality constraints

$g(x,y) \leq 0$ and $-g(x,y) \leq 0$. However, our main result will require the constraining function to be convex. Hence the trick works only if $g(x,y)$ is linear.

Example 2 A consumer wants to maximize the utility function $u(x,y)$ subject to the budget constraint

$$p_x x + p_y y \leq b \qquad \text{or} \qquad b - p_x x + p_y y \geq 0$$

and the non-negativity constraints

$$\begin{aligned} x &\geq 0 \\ y &\geq 0 \end{aligned}$$

A pair of numbers (x,y) is *admissible* for the optimization problem if it is in the domain D and if it satisfies all constraints, i.e. if

$$g_k(x,y) \geq 0 \qquad \text{for all} \qquad 1 \leq k \leq K$$

A *solution* to the constrained maximization problem is an admissible pair of numbers (x^*,y^*) such that

$$f(x^*,y^*) \geq f(x,y)$$

holds with respect to all admissible pairs (x,y).

Let (x^*,y^*) be a solution to the constrained maximization problem. Then we say that the k-th constraint g_k (whereby $1 \leq k \leq K$) is *binding at the optimum* (x^*,y^*) if $g_k(x^*,y^*) = 0$. If $g_k(x^*,y^*) > 0$ we say that it is *not binding.*

If the k-the constraint is not binding, then it is superfluous in the sense that we could leave it out from the constrained optimization problem and the solution would still be a solution.

If the constrained is binding then we expect the solution in general to change if the constraint is left out. If this is indeed the case, we say the constraint is *strictly binding.* otherwise it is *just binding.*

1.1 The Lagrangian Approach (Kuhn-Tucker)

Idea: We transform the constrained optimization problem into an unconstrained optimization problem (to be solved via the usual first-order conditions) and a pricing problem.

Instead of maximizing the objective function we maximize the so-called *Lagrangian* $\mathcal{L}$. In the Lagrangian we add to the objective function a weighted sum of the constraining functions. The weights on the constraining functions are called the *Lagrangian multipliers. They must be non-negative.*

Thus the Lagrangian takes the form

$$\mathcal{L}(x,y) = f(x,y) + \lambda_1 g_1(x,y) + \lambda_2 g_2(x,y) + \ldots + \lambda_K g_K(x,y)$$

The Lagrangian multipliers $\lambda_1, \lambda_2, \ldots \lambda_K$ should be understood as *shadow prices* (meaning fictitious prices which lurk behind the problem) for the constraints. The amount λ_k gets subtracted from the objective function for each unit of violation of constraint. Foe instance, to violate constraint g_1 five units ($g_1(x,y) = -5$) "costs" $-5\lambda_1$, i.e. $-5\lambda_1$ get subtracted from the Lagrangian. Thus if (x, y) violates constraint k and hence $g_k(x, y) < 0$ we subtract the negative amount

$$\lambda_k g_k(x, y)$$

from the objective function. When maximizing the Lagrangian one hence tries to get the value of the objective function large and the violations of the constraints small.

The *pricing problem* is the problem to get the shadow prices just right such that the solution of the unconstrained problem coincides with the original constrained optimization problem.

The following statement shows that this approach makes sense.

Theorem 1 *Suppose*

- *there are non-negative Lagrangian multipliers* $\lambda_1, \lambda_2, \ldots \lambda_K$
- *and an unconstrained maximum* (x^*, y^*) *for the Lagrangian (formed using these parameters)*
- *which is admissible for the constrained optimization problem*
- *and where the Lagrange multiplier* λ_k *is zero whenever the constraint* g_k *does not bind at* (x^*, y^*).

Then (x^*, y^*) *is a solution to the original constrained optimization problem.*

Proof. Notice first for *any* admissible pair (x, y) that

$$f(x, y) \leq \mathcal{L}(x, y) \tag{1}$$

because $\lambda_k g_k(x, y) \geq 0$ for all k since (x, y) does not violate any constraints and all λ_k are non-negative by assumption.

Now consider the pair (x^*, y^*) as described in the theorem. Because (x^*, y^*) is admissible we have $g_k(x^*, y^*) \leq 0$ for every k. If constraint k binds at (x^*, y^*) then $\lambda_k g_k(x^*, y^*) = 0$. By assumption $\lambda_k = 0$ holds whenever $g_k(x^*, y^*) > 0$. Hence the *complementarity conditions*

$$\lambda_k g_k(x^*, y^*) = 0 \tag{2}$$

hold for *every* k. But this means that

$$\mathcal{L}(x^*, y^*) = f(x^*, y^*) \tag{3}$$

Since (x^*, y^*) is an unconstrained maximum of the Lagrangian we have

$$\mathcal{L}(x, y) \leq \mathcal{L}(x^*, y^*) \tag{4}$$

for all pairs (x, y) in the domain, in particular for all admissible pairs (x, y).

Combining (1), (4) and (3) we have

$$f(x, y) \leq f(x^*, y^*)$$

as was to be shown. ■

Under convexity conditions one can use an important existence theorem, the Hahn-Banach separation theorem, to show that for every solution to the constrained optimization problem (x^*, y^*) there exists non-negative Lagrange multipliers such that the complementarity conditions hold and such that (x^*, y^*) maximizes the Lagrangian. We may return to this point.

2 A first example

We illustrate in a simple example how the theorem can be used to find the optimum.

Example 3 Let S be the square consisting of points (x, y) with $-1 \leq x, y \leq 1$. Use the Lagrangian method to find for a given point $P = (a, b)$ find, the point $Q = (x^*, y^*)$ in the square with the smallest distance to P.

Solution 1 By the law of Pythagoras the distance between (x, y) and (a, b) is

$$r = \sqrt{(x-a)^2 + (y-b)^2}$$

Our objective is to minimize r subject to the constraint that (x, y) is in C. This is the same as minimizing r^2 or maximizing $-r^2$ under this constraint. Hence our optimization problem is to maximize the objective function

$$z = f(x, y) = -(x-a)^2 - (y-b)^2$$

subject to the four constraints $-1 \leq x$, $x \leq 1$, $-1 \leq y$ and $y \leq 1$ or, alternatively,

$$g_1(x, y) = 1 + x \geq 0 \quad (5)$$

$$g_2(x, y) = 1 - x \geq 0 \quad (6)$$

$$g_3(x, y) = 1 + y \geq 0 \quad (7)$$

$$g_4(x, y) = 1 - y \geq 0 \quad (8)$$

The Lagrangean is

$$\begin{aligned} \mathcal{L}(x, y) &= -(x-a)^2 - (y-b)^2 \\ &\quad + \lambda_1(1+x) + \lambda_2(1-x) + \lambda_3(1+y) + \lambda_4(1-y) \end{aligned}$$

and has the partial derivatives

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial x} &= -2(x-a) + \lambda_1 - \lambda_2 \\ \frac{\partial \mathcal{L}}{\partial x} &= -2(y-b) + \lambda_3 - \lambda_4 \end{aligned}$$

Notice that the Lagrangian is always strictly concave and has hence a unique maximum.

The optimum can be in the interior of the square or on the boundary. If it is on the boundary it can be on any of the four sides or on any of the four corner points. Hence there are nine possibilities to be studied. In each case we try to satisfy the assumptions of the above theorem.

Case (1) Suppose the optimum is in the interior of the square: $-1 < x^* < 1$, $-1 < y^* < 1$, i.e. none of the four constraints is binding, all four constraints are satisfied with a strict inequality sign '>'. When no constraint is binding the complementary condition requires that all Lagrange multipliers are zero, so the Lagrangian is just the objective function in this case. The first order conditions are the same as for an unconstrained optimum.

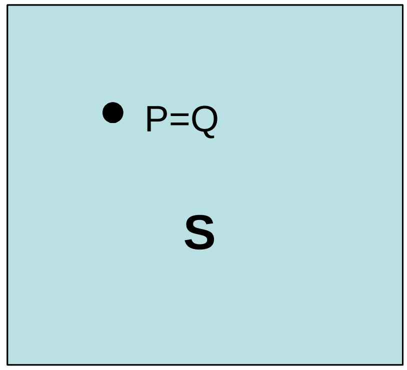

In this case the optimum has to be a stationary point of the objective function. This gives the conditions

$$\begin{aligned} \frac{\partial f}{\partial x} &= -2(x-a) = 0 \\ \frac{\partial f}{\partial y} &= -2(y-b) = 0 \end{aligned}$$

There is only one solution to this system of equations: $(x^*, y^*) = (a, b)$. However, this solution can be a solution to our constrained optimization problem only if all four constraints are satisfied. Moreover, we have assumed that they are satisfied strictly. So, there can be no interior solution unless $-1 < a < 1$, $-1 < b < 1$ is satisfied, in which case (a, b) is our only candidate for an optimum.

Interpretation: The only point with distance 0 to (a, b) is (a, b) itself. So, if (a, b) is in the square it is clearly the solution to our optimization problem. If (a, b) is outside the square it cannot be a solution. Although the above analysis assumed the solution to be in the interior, the result still makes sense if it is on the boundary.

Case (2) Suppose the optimum is on the upper side of the square, but not a corner point, i.e. constraint (8) is binding, the other three constraints hold with strict inequality.

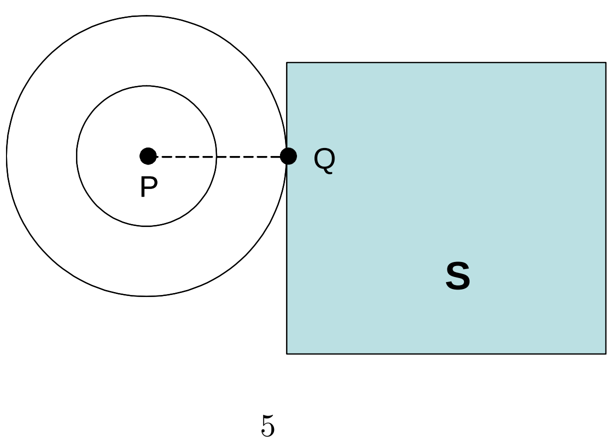

By the complementarity conditions all Lagrange multipliers except λ_4 must be zero. So the Lagrangian is

$$\mathcal{L}(x,y) = f(x,y) + \lambda_4 (1-y)$$

So the first order conditions are

$$\begin{aligned}\frac{\partial \mathcal{L}}{\partial x} &= -2(x-a) = 0 \\ \frac{\partial \mathcal{L}}{\partial y} &= -2(y-b) - \lambda_4 = 0\end{aligned}$$

to which we add the assumption that constraint (8) is binding, i.e.

$$y = 1.$$

The three equations have the unique solution

$$x^* = a, y^* = 1, \lambda_4 = -2(1-b).$$

For $(x^*, 1) = (a, 1)$ to satisfy all constraints we must have $-1 \leq a \leq 1$. In addition, the Lagrangian method requires λ to be nonnegative, which will be the case only if $1 - b \leq 0$ or $b \geq 0$.

Cases (3, 4, 5) The optimum is on one of the other three sides. The analysis is symmetrical.

Case (6) Suppose the optimum is the corner point $(1, 1)$. Then constraints (6) and (8) are binding while (5) and (7) are not.

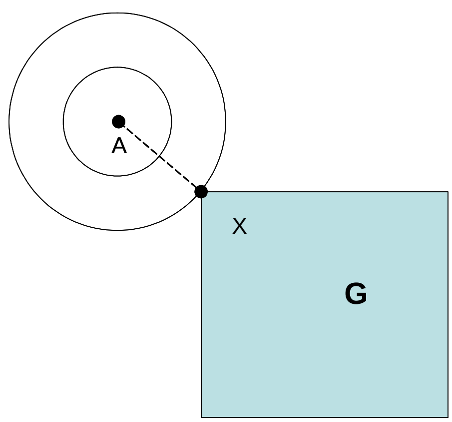

By the complementarity conditions we must therefore have $\lambda_1 = \lambda_3$ and the Lagrangian takes the form

$$\mathcal{L}(x,y) = f(x,y) + \lambda_2 (1-x) + \lambda_4 (1-y)$$

A stationary point must satisfy

$$\begin{aligned}\frac{\partial \mathcal{L}}{\partial x} &= -2(x-a) - \lambda_2 = 0 \\ \frac{\partial \mathcal{L}}{\partial y} &= -2(y-b) - \lambda_4 = 0\end{aligned}$$

to which we add the conditions

$$x = 1$$
$$y = 1.$$

The unique solution to this four equations is $(x^*, y^*) = (1, 1)$, $\lambda_2 = -2(1-a)$, $\lambda_4 = -2(1-b) = 0$. To have non-negative Lagrange multipliers we need

$$a \geq 1 \quad \text{and} \quad b \geq 1.$$

Cases (7, 8, 9) These cases are treated symmetrically. It is recommended to do at least one.

The following table describes for each pair (a, b) what the optimum (x^*, y^*) is

	$b \leq -1$	$-1 < b < 1$	$1 \leq b$
$a \leq -1$	$(-1,-1)$	$(-1, b)$	$(-1, 1)$
$-1 < a < 1$	$(a, -1)$	(a, b)	$(a, 1)$
$1 \leq a$	$(1, -1)$	$(1, b)$	$(1, 1)$

Exercise 1 Let Δ be the triangle spanned by the points $(-1, 0)$, $(1, 1)$ and $(1, -1)$. Sketch the triangle. For each point (a, b) in the plane find the point (x^*, y^*) which is closest to it.

3 Utility maximization

Example: Maximize $u(x, y) = (x+1)(y+1)$, subject to the budget constraint $p_x x + p_y y \leq b$ and the non-negativity constraints $x, y \geq 0$.

Where can the consumer optimum be in the budget set carved out by the budget- and the non-negativity constraints?

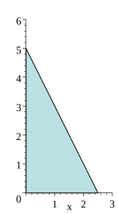

In principle there are 7 different possibilities to consider: It could be in the interior of the triangle, on one of the three sides or it could be one of the three corner points.

$u(x, y) = (x+1)(y+1)$ is "monotonic", more of each commodity is better.

Lemma 1 *The budget constraint must be binding in the consumer optimum, i.e. the optimum is on the budget line.*

Proof. If (x, y) is a consumption bundle which costs less than the budget then one can, for instance, slightly increase the consumption of x without violating the budget constraint. The first factor in $(x+1)(y+1)$ increases and hence, since $y+1>0$ the whole product. Utility goes up and so (x, y) cannot be the optimum. ■

This argument holds generally for all *monotonic preferences*.

This reduces the search to three possibilities: The budget line and its two corner points. (Different cases: "bads", satiated consumers.)

Three possibilities:

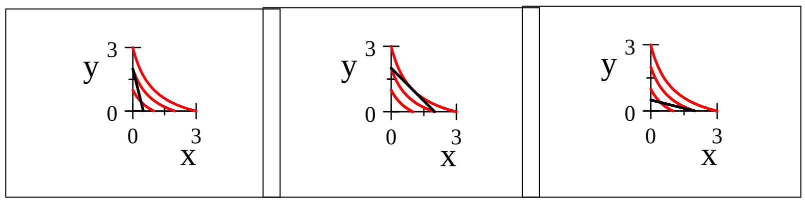

Optimum (0,2) Optimum (1,1) Optimum (2,0)

For a utility function like $u(x, y) = xy$ *one can say even more*, namely that consumption of both commodities must be strictly positive in optimum. (Utility is zero when $x = 0$ or $y = 0$ while a strictly positive utility can be achieved with a strictly positive budget.) Thus *only* the budget constraint can be binding.

The Lagrangian:

$$\begin{aligned}\mathcal{L}(x, y) &= u(x, y) + \lambda_1 (b - p_x x - p_y y) + \lambda_2 x + \lambda_3 y \\ &= (x+1)(y+1) + \lambda_1 (b - p_x x - p_y y) + \lambda_2 x + \lambda_3 y\end{aligned}$$

where $p_x, p_y, b > 0$

FOC:

$$\begin{aligned}\frac{\partial \mathcal{L}}{\partial x} &= \frac{\partial u}{\partial x} - \lambda_1 p_x + \lambda_2 = (y+1) - \lambda_1 p_x + \lambda_2 = 0 \\ \frac{\partial \mathcal{L}}{\partial y} &= \frac{\partial u}{\partial y} - \lambda_1 p_y + \lambda_3 = (x+1) - \lambda_1 p_y + \lambda_3 = 0\end{aligned}$$

3.1 Case 1: Only the budget equation binds

Hence $\lambda_2 = \lambda_3 = 0$ by the complementarity condition. We get the FOC

$$\begin{aligned}\frac{\partial \mathcal{L}}{\partial x} &= \frac{\partial u}{\partial x} - \lambda_1 p_x = (y+1) - \lambda_1 p_x = 0 \\ \frac{\partial \mathcal{L}}{\partial y} &= \frac{\partial u}{\partial y} - \lambda_1 p_y = (x+1) - \lambda_1 p_y = 0\end{aligned}$$

and the budget constraint must hold with equality

$$p_x x + p_y y = b.$$

Notice that if x and y are postive, then the FOC imply $\lambda_1 > 0$.

When only one constraint binds, it is easy to eliminate the Lagrange multiplier:

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial x} &= \frac{\partial u}{\partial x} - \lambda_1 p_x = 0 \\ \frac{\partial \mathcal{L}}{\partial y} &= \frac{\partial u}{\partial y} - \lambda_1 p_y = 0 \end{aligned}$$

implies

$$\frac{\partial u}{\partial x} \Big/ \frac{\partial u}{\partial y} = \frac{p_x}{p_y}$$

Together with

$$p_x x + p_y y = b.$$

we must then solve a system of two equations with two unknowns.

The FOC imply

$$\frac{\partial u}{\partial x} = \lambda_1 p_x \qquad \frac{\partial u}{\partial y} = \lambda_1 p_y$$

Division of the two left hand sides and the two right hand sides yields

$$\frac{\partial u}{\partial x} \Big/ \frac{\partial u}{\partial y} = \frac{p_x}{p_y}$$

Thus the marginal rate of substitution (see previous lectures) must equal the price ratio or, in other words, in the consumer optimum the indifference curve is tangential to the budget line.

In our particular example this yields

$$\frac{y+1}{x+1} = \frac{p_x}{p_y} \qquad y = \frac{p_x}{p_y}(x+1) - 1$$

Substitution into the budget equation yields.

$$\begin{aligned} b &= p_x x + p_y y = p_x x + p_y \left(\frac{p_x}{p_y}(x+1) - 1 \right) \\ &= 2 p_x x + p_x - p_y \end{aligned}$$

and so

$$x^* = \frac{b - p_x + p_y}{2 p_x} \qquad y^* = \frac{b + p_x - p_y}{2 p_y}$$

where the last formula holds because

$$y = \frac{p_x}{p_y}(x+1) - 1 = \frac{p_x}{p_y} \frac{b - p_x + p_y}{2 p_x} + \frac{2 p_x}{2 p_y} - \frac{2 p_y}{2 p_y} = \frac{b + p_x - p_y}{2 p_y}$$

$$x^* = \frac{b - p_x + p_y}{2 p_x} \qquad y^* = \frac{b + p_x - p_y}{2 p_y}$$

can only be the solution when both numbers are non-negative. This requires

$$\begin{aligned} b - p_x + p_y &\geq 0 \qquad b + p_x - p_y \geq 0 \\ b &\geq p_x - p_y \qquad b \geq -(p_x - p_y) \\ b &\geq |p_x - p_y| \end{aligned}$$

where || denotes the "absolute value"

$$|a| = \begin{cases} a & \text{if } a \geq 0 \\ -a & \text{if } a < 0 \end{cases}$$

Intuitively, the price difference cannot be too large in comparison to the budget.

To summarize, provided $b \geq |p_x - p_y|$ the Lagrangian approach yields a positive Lagrange multiplier λ_1 (see the argument further above) and the non-negative solution

$$x^* = \frac{b - p_x + p_y}{2p_x} \qquad y^* = \frac{b + p_x - p_y}{2p_y}$$

Provided we can show that this solution is indeed an unconstrained maximum of the Lagrangian (and not a minimum etc.) it is the solution to our constrained optimization problem.

3.2 Case 2: The budget equation binds and x=0

So one of the non-negativity constraint is binding. Thus the first order conditions

$$\begin{aligned} \frac{\partial \mathcal{L}}{\partial x} &= \frac{\partial u}{\partial x} - \lambda_1 p_x + \lambda_2 = (y+1) - \lambda_1 p_x + \lambda_2 = 0 \\ \frac{\partial \mathcal{L}}{\partial y} &= \frac{\partial u}{\partial y} - \lambda_1 p_y + \lambda_3 = (x+1) - \lambda_1 p_y + \lambda_3 = 0 \end{aligned}$$

must holds together with the budget equation

$$p_x x + p_y y = b$$

and

$$x = 0$$

The budget equation simplifies to $p_y y = b$ and our only solution candidate is

$$x^* = 0 \qquad y^* = b/p_y$$

Because $y^* > 0$ the non-negativity constraint $y = 0$ does not bind and therefore $\lambda_3 = 0$ by the complementarity conditions. The FOC simplify to

$$\begin{aligned} b/p_y + 1 - \lambda_1 p_x + \lambda_2 &= 0 \\ 1 - \lambda_1 p_y &= 0 \end{aligned}$$

So $\lambda_1 = 1/p_y \geq 0$ and the first FOC yields

$$\begin{aligned} b/p_y + 1 - \frac{p_x}{p_y} + \lambda_2 &= -b/p_y - 1 + \frac{p_x}{p_y} \\ \frac{b + p_y - p_x}{p_y} + \lambda_2 &= 0 \\ \lambda_2 &\geq \frac{p_x - p_y - b}{p_y} \end{aligned}$$

For λ_2 to be non-negative we need that the price difference $p_x - p_y$ is bigger than the budget.

Case 3: The budget equation binds and y=0

This case is handled completely symmetrically to case 2.

3.3 Summary

Apart form showing that we have indeed found unconstrained optima of the Lagrangian we get the following result

Theorem 2

- *When the price of x is very high, namely when $p_x \geq p_y + b$ the consumer only wants to buy y and so $x^* = 0$, $y^* = b/p_y$.*
- *When the price of y is very high, namely when $p_y \geq p_x + b$ the consumer only wants to buy y and so $x^* = b/p_x$, $y^* = 0$.*
- *In all other cases the consumer wants to buy of both commodities the amounts*

$$x^* = \frac{b - p_x + p_y}{2p_x} \qquad y^* = \frac{b + p_x - p_y}{2p_y}$$