

We consider the situation when we have a dataset describing a number of persons each of whom is characterized by two quantities  $x$  and  $y$ . Both  $x$  and  $y$  are measured before and after a certain treatment. There are two treatments, A and B, and each person is subjected to just one of them. Our goal is to compare how A and B affect both  $x$  and  $y$ . More precisely: we want to find support for a statement of the type 'A affects  $x$  more than does B while B affects  $y$  more than does A'. Several methods that easily come to mind will be described and I shall try to explain why they are all inappropriate. Then I shall describe a new method which perhaps at first sight may seem a bit strange. As illustration will be used data from a clinical trial of blood-pressure lowering drugs.