

Models for the Assessment of Treatment Improvement: The Ideal and the Feasible

P. C. Álvarez-Esteban, E. del Barrio, J. A. Cuesta-Albertos and C. Matrán

Abstract. Comparisons of different treatments or production processes are the goals of a significant fraction of applied research. Unsurprisingly, two-sample problems play a main role in statistics through natural questions such as “Is the the new treatment significantly better than the old?” However, this is only partially answered by some of the usual statistical tools for this task. More importantly, often practitioners are not aware of the real meaning behind these statistical procedures. We analyze these troubles from the point of view of the order between distributions, the stochastic order, showing evidence of the limitations of the usual approaches, paying special attention to the classical comparison of means under the normal model. We discuss the unfeasibility of statistically proving stochastic dominance, but show that it is possible, instead, to gather statistical evidence to conclude that slightly relaxed versions of stochastic dominance hold.

Key words and phrases: Stochastic dominance, similarity, two-sample comparison, trimmed distributions, winsorized distributions, Behrens–Fisher problem, index of stochastic dominance.

1. INTRODUCTION

Comparison is an essential activity in any field of life, one upon which a significant part of human knowledge is founded. In fact, one of the main achievements of mankind—numbers—are just a wonderful sophistication of the comparison process. Whether by curiosity or necessity we are continuously involved in comparing objects, leading to assessments like bigger/smaller,

shorter/taller, better/worse and so forth. It is therefore natural the prominent role played in statistics by procedures looking for some kind of ordering. In fact, “Two-Sample Problems”, that is, comparing two populations or two treatments, is probably the most common situation encountered in statistical practice. Not surprisingly, every textbook on statistics explores the topic to some extent.

In many cases, the practitioner using a two-sample procedure has the goal of gathering evidence to conclude that a new treatment is better than the old standard. To fix ideas, let us assume that treatment refers to a particular training program for athletes. To assess the possible improvement provided by a new training program, the researcher collects some experimental data from athletes training under the two different programs. Of course, not everything comes from the type of training and one expects that a naturally talented athlete will perform better, whatever the training program, than a less gifted one. In the simplest case in which performance is measured in terms of a simple univariate outcome, we can think of a training program as a nondecreasing transform of the level of natural talent. If, in some scale, this talent level is measured as T and the different training programs result in $h_{\text{old}}(T)$

P. C. Álvarez-Esteban is Professor, Departamento de Estadística e Investigación Operativa and IMUVA, Universidad de Valladolid, Paseo de Belén, 7. 47011-Valladolid, Spain (e-mail: pedroc@eio.uva.es). E. del Barrio is Professor, Departamento de Estadística e Investigación Operativa and IMUVA, Universidad de Valladolid, Paseo de Belén, 7. 47011-Valladolid, Spain (e-mail: tasio@eio.uva.es). J. A. Cuesta-Albertos is Professor, Departamento de Matemáticas, Estadística y Computación, Universidad de Cantabria, Avda. de los Castros, 48. 39005-Santander, Spain (e-mail: cuestaj@unican.es). C. Matrán is Professor, Departamento de Estadística e Investigación Operativa and IMUVA, Universidad de Valladolid, Paseo de Belén, 7. 47011-Valladolid, Spain (e-mail: carlos.matran@uva.es).

and $h_{\text{new}}(T)$ levels of performance, respectively, we would say that the new treatment is better than the old if $h_{\text{old}}(t) \leq h_{\text{new}}(t)$ for all t .

What type of conclusion is drawn from the most standard use of the two-sample procedures? The most commonly used test, namely, the t -test, even in the Welch version related to the famous Behrens–Fisher problem, would simply aim at rejecting that the mean of $h_{\text{old}}(T)$ is greater than the mean of $h_{\text{new}}(T)$. However, as we show in this paper, even under the normality assumptions implicit in the use of the t -test, evidence of a significantly greater mean under the new treatment is compatible with a worse performance of, say, 40% of athletes. We believe that practitioners should be aware of this fact. We argue in this paper that, most often, they would rather be interested in gathering evidence for stochastic dominance than for an increase in mean values.

In this goal of gathering evidence to support the claim that the new treatment yields an improvement over the old one, we cannot forget that testing hypothesis theory is designed to provide evidence to reject the null hypothesis and that lack of rejection does not mean evidence for the null. While this is a well-known fact in the statistical community, in the absence of other approaches, practitioners often resort to widely used procedures without full conscience of their true meaning. Perhaps the best example of such a situation is the generalized use of goodness of fit tests, as the Kolmogorov–Smirnov test, as a way to justify a parametric model assumption, such as normality. A closer example to our present framework is that of testing homogeneity in a two sample setup. In both situations, regardless of the obtained result, we will not be able to confirm the model but, at best, we would just get lack of statistical evidence to reject it. This fact has been pointed out by several authors, notably by Dette and Munk (1998) or Munk and Czado (1998). In the particular case of stochastic dominance, we should test the null that stochastic dominance does not hold against the alternative that it does hold if we want evidence supporting it (hence, evidence supporting that the new treatment is better). Unfortunately, no reasonable statistical test can help in this task; see Berger (1988) and our discussion in Section 2.1 below.

On the other hand, a model is merely an approximation to reality, so, in order to validate a model, we should be conscious of what are the admissible deviations to the model. This is the starting point for the discussion of practical versus statistical significance in Hodges and Lehmann (1954) continued in a series

of papers (see, e.g., Rudas, Clogg and Lindsay, 1994, Liu and Lindsay, 2009, Álvarez-Esteban et al. 2008, 2012, 2014) having the common goal of testing the approximate validity of statistical hypothesis. In this paper, we discuss two relaxations of the stochastic dominance model for which there are consistent statistical tests. More precisely, we show that there are consistent tests that allow the practitioner to conclude, up to some small probability of error, that the new treatment is within a small neighborhood of being better than the old one.

The remaining sections of this paper are organized as follows. In Section 2, we further discuss the convenience of considering stochastic order rather than growth in mean when trying to assess the improvement given by a new treatment. Details about the above mentioned lack of valid inferential methods for concluding stochastic dominance are also given, together with two relaxations of stochastic dominance, which can be used to produce indices of deviation from the ideal model of stochastic order. We include a subsection that illustrates the behavior of these indices through the important example of distributions that differ only in changes in location or scale, showing that growth in the mean is well compatible with a worse performance under a new treatment for a very substantial fraction of the population. We invite to a careful inspection of the graphics in Figures 3 and 5 to get a visual impact of the departures of stochastic dominance measured through such indices, when compared with changes in location and scale in the normal model. Section 3 provides valid inferential methods for gathering statistical evidence that the relaxed stochastic order models hold; hence, showing that these deviations from the ideal model of stochastic dominance are tractable models from the point of view of statistical inference. We also provide a simulation study showing the performance of the inferential methods in finite samples. Finally, the proofs of some results introduced in this work are included in an Appendix.

2. MODELS OF TREATMENT IMPROVEMENT

2.1 Distributional Dominance Versus Mean Comparisons

Let us briefly explore the use and real meaning of the most common approach to assess improvement in two sample problems. Assuming independence between the samples and normality in the parent distributions, the t -test is based on the comparison of the means of the distributions. However, relations between

the means, or any other feature of the distributions, must be cautiously evaluated to assess some kind of improvement in a production process or of a treatment with respect to another.

To motivate our discussion, let us assume that the probability laws of the variable of interest under two different production processes are normal, say $P_i = N(\mu_i, \sigma_i^2), i = 1, 2$. If we could conclude that $\mu_1 > \mu_2$ we would be only allowed to claim that “in the mean” the first process produces larger values than the second. To better explain the meaning of such a statement, we can resort to the strong law of large numbers: for large enough samples obtained from both processes, the mean of the sample obtained from P_1 would be almost surely greater than that of the obtained from P_2 . We should stress the fact that this statement does not depend on the values σ_i . Thus it is compatible with the situations displayed in Figure 1. The relevant question is whether these situations are compatible with our intuitive understanding of the statement that the first process leads to greater values than the second.

An informal statement such as “men are taller than women” can be better detailed saying that a short man would not be as short among women, a medium-sized man would be tall among women and a tall man would be even taller among women. These comparisons involve in a natural way the relative position or status

of every item in both populations. With greater precision, the whole comparison involves checking the new status of each element of the first population when we consider it as an element of the second population. If it is always greater, then we could say that the variable in the first population is greater than in the second.

To analyze if this relation holds, a notable simplification is achieved through a previous arrangement of each population ordering their elements by their status. In that way, the relative position of each element in its population is easily obtained, and normalization in each population by its size to get comparable status leads to a simple relation. If F and G are respectively the distribution functions of the variable of interest in the first and the second populations, and F^{-1}, G^{-1} are the corresponding quantile functions,

$$(1) \quad F^{-1}(t) \leq G^{-1}(t) \quad \text{for every } t \in (0, 1),$$

would mean that the variable in the first population is lower than in the second. We recall that for a general distribution function on the real line, F , the associated quantile function, that we will denote as F^{-1} , is defined as

$$F^{-1}(t) := \min\{x : t \leq F(x)\}, \quad t \in (0, 1).$$

It is well known that the relation (1) is equivalent to the classical definition of stochastic order, usually

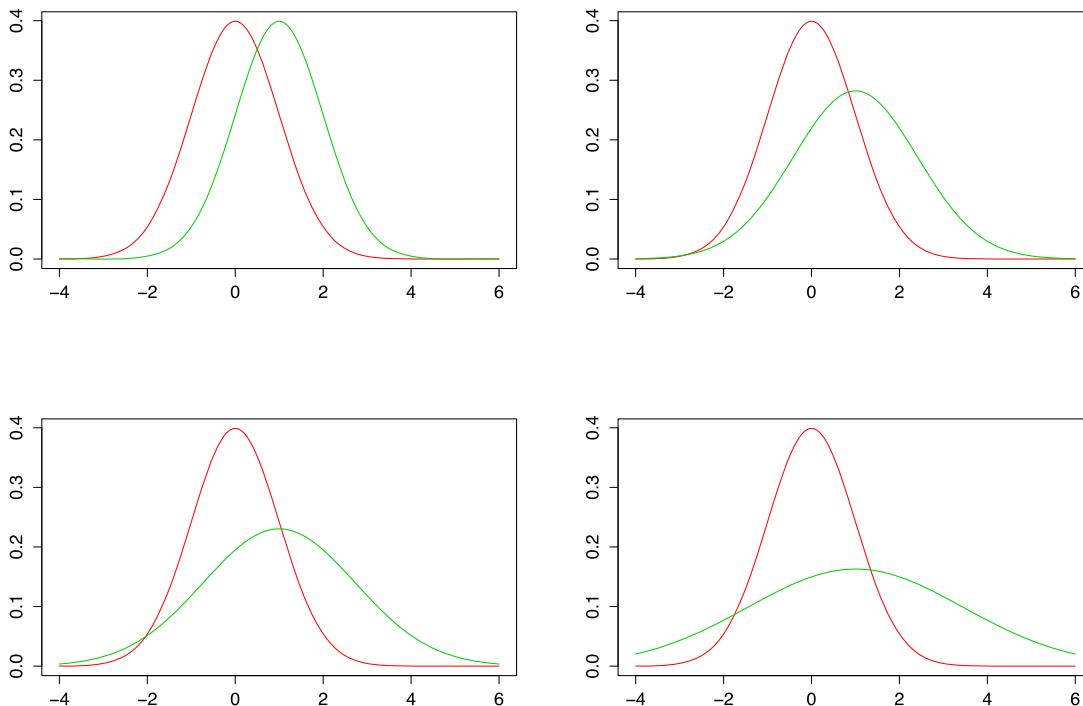


FIG. 1. Red: standard normal density; green: densities of $N(1, \sigma^2), \sigma^2 = 1, 2, 3, 6$.

attributed to Lehmann (1955), but already used at least in Mann and Whitney (1947). We say that F is stochastically smaller than G , and write $F \leq_{st} G$, if

$$(2) \quad F(x) \geq G(x) \quad \text{for every } x \in \mathbb{R}.$$

In the econometric literature, where more general classes of stochastic orders are considered, usually linked to preferences related with families of utility functions, this relation is often invoked as first-order stochastic dominance (see, e.g., Shaked and Shanthikumar, 2007 or Müller and Stoyan, 2002 for several extensions of the concept).

An interesting fact about quantile functions is that if T is uniformly distributed on $(0, 1)$ then $F^{-1}(T)$ has distribution function F . Let us return for a moment to the discussion in the Introduction and think of F and G as the distribution functions of the performances of athletes training under the old and new programs, respectively. Let T be a measure of the natural talent of a randomly chosen athlete and let F_T be its d.f. If we assume that F_T is continuous, then, it is well known that $T^* := F_T(T)$ is uniform on $(0, 1)$. Therefore, just making a modification on the measurement scale, we have that the variable giving the natural talent of the athletes is uniform on $(0, 1)$. Then we can see $F^{-1}(T^*)$ and $G^{-1}(T^*)$ as the effects of the training programs on that natural talent. Hence, F^{-1} plays the role of h_{old} and G^{-1} that of h_{new} in the discussion

in the Introduction and we see from the interpretation there that the new training program was better than the old if $h_{old}(t) \leq h_{new}(t)$ for all t coincides with first order stochastic dominance.

In view of the arguments above, we think that a sound answer to question “is the new treatment better than the old” should be based on the assessment of stochastic order. A look at Figure 2 shows that this cannot be done by simply comparing the means. In fact, the mean is the same for all distributions in green, but stochastic order only holds in the comparisons $N(0, 1)$ versus $N(1, 1)$. Specific inferential methods for assessing $F \leq_{st} G$ are needed.

There is an abundant literature in statistical and econometric journals concerning testing problems related to stochastic dominance. Some references, tracing back to Mann and Whitney (1947), belong to an order restricted inferential approach, that is, assuming that stochastic order holds they focus on concluding that *strict* stochastic order holds [$F <_{st} G$ if $F(x) \geq G(x)$ for all x with $F(x) > G(x)$ for at least one x]. More precisely, they consider the problem of testing the null hypothesis $H_0 : F = G$ against the alternative $H_a : F <_{st} G$. Obvious as it may be, it is relevant to note that some caution should be adopted in applying these procedures, since, both H_0 and H_a can be simultaneously false.




FIG. 2. Red: standard normal quantile function; green: quantile functions of $N(1, \sigma^2)$, $\sigma^2 = 1, 2, 3, 6$.

A different testing problem with a number of references in the literature (see, e.g., McFadden, 1989, Anderson, 1996, Barrett and Donald, 2003, Davidson and Duclos, 2000, Linton, Maasoumi and Whang 2005, Linton, Song and Whang, 2010) is that of testing the null $H_0: F \leq_{st} G$, versus $H_a: F \not\leq_{st} G$. This is a kind of goodness-of-fit test. The statistical meaning of not rejecting the null is simply to acknowledge that there is not evidence enough to guarantee that stochastic dominance does not hold. However, this “accepting” the null is sometimes invoked as a guarantee that one random variable is stochastically larger than other, but this is simply wrong.

The available analyses do not address the main goal of gathering statistical evidence to assess that stochastic dominance, $F \leq_{st} G$, holds. This would be the natural result of rejecting the null in the problem of testing

$$(3) \quad H_0: F \not\leq_{st} G, \quad \text{versus } H_a: F \leq_{st} G.$$

Unfortunately, as already noted in Berger (1988), Davidson and Duclos (2013) and Álvarez-Esteban et al. (2014), the statistical assessment of distributional dominance is impossible because small variations in the tails of a distribution could avoid or facilitate a relation of stochastic dominance. In fact, there is no good α -level test for (3): the “no data” test, rejecting H_0 with probability α regardless of the data is uniformly most powerful. This is shown in Berger (1988) in the one-sample setting, but the result can be easily generalized to the two-sample setup considered here.

2.2 Relaxations of Stochastic Dominance

As it often happens in statistics, the concept of stochastic dominance is excessively rigid as to try to confirm it on the basis of a sample. It is a too strong assumption in problems in which one is inclined to believe that a population X is somehow smaller than another population Y . This difficulty seems to be a big justification for the common practice of basing the comparisons on features of the distributions, like the means, when trying to assess some kind of order between distributions. For some parametric models, as in the normal model, this approach has the additional advantage of leading to true stochastic order for distributions with the same variance. Since optimal testing for this problem can be achieved through an exact test, the two-sample t -test, the approach seems almost perfect. For the celebrated Behrens–Fisher problem, when the variances are not assumed equal, approximations such as Welch’s proposal give satisfactory enough solutions to face the comparison of means problem. However,

even in the normal model, this may be giving a right answer to a wrong question.

Stochastic order is a 0–1 relation. It is either true or false (of course, the same can be said for higher order choices of stochastic dominance). In the case of normal laws, for instance, stochastic order holds only in the case of equal variances and increasing means; see Section 2.3 for related results on location-scale models (LS-models in the sequel). However, looking back at Figure 2, it is tempting to say that the degree of deviation from stochastic order is higher in the example in the lower right corner than in those in the lower left or upper right corners. We could say that in these last cases stochastic dominance nearly holds or, even, that “in practice”, it holds. Some measurement of the level of agreement with stochastic order would be helpful.

Motivated by the unfeasibility of consistently testing (3), Berger (1988) considered the idea of “restricted stochastic dominance”, which amounts to looking for the relation $F(x) \geq G(x)$ on a fixed closed interval, excluding the tails of the sampled distribution. The choice of the interval is somewhat arbitrary. The same approach had already been considered in Lehmann and Rojo (1992), as a weak version of the stochastic order, stressing the fact that it is not always appropriate to require that the comparison holds for all values of x . An analogous version has been developed in the two-sample setting by Davidson and Duclos (2013). Given the equivalence between (1) and (2), it would make sense, as well, to fix an interval contained in $(0, 1)$ and check whether (1) holds within this interval. On the other hand, the normalization given by the quantile transform gives some additional advantage. Rather than facing the arbitrary choice of the interval on which we want to check that (1) holds, we can look at the length of the the set where it does not hold, namely,

$$(4) \quad \gamma(F, G) = \ell(t \in (0, 1) : F^{-1}(t) > G^{-1}(t)),$$

where ℓ stands for Lebesgue measure. This yields a useful index to measure how far F and G are from stochastic order, with $\gamma(F, G) = 0$ corresponding to perfect fit. If we turn back to the example of the old and new training programs for an athlete, then $\gamma(F, G) = 0.05$ means that 95% of athletes get better results with the training program associated to G than with that of F . When restricted to a specific model as the normal model (and more generally to LS-models) the computation and the meaning of this index is easy and very informative. The contour-plot in Figure 3 gives a nice insight into the fact that moderate and even high levels of disagreement with stochastic order [up




FIG. 3. Contour-plot of $\gamma(N(0, 1), N(\mu, \sigma^2))$ as in (4) for different values of μ (X-axis) and σ (Y-axis).

to $\gamma(F, G) \sim 0.5]$ are compatible with an increase in mean, even in the normal model. As an example, if F denotes the standard normal law, $N(0, 1)$, and G corresponds to the law, $N(0.337, 1.5^2)$, then we have $\gamma(F, G) = 0.25$. Again, in the training program example, we see that the new program can yield an improvement in the mean performance of athletes and, yet, result in worse results for 25% of them.

We recall that our initial motivation was to discuss the suitability of the usual methods involved in the validation of domination or improvement. From Figure 3 and the above discussion, the inadequacy of the two sample t -test to validate a real improvement for most of the population (unless both distributions satisfy some strong additional assumptions such as normality plus equal variances) should be obvious. Of course, this is not an objection to the use of, say, the Welch version of the t -test for testing an increase in mean. The key point is the meaning of these mean comparisons for the task of showing improvement of treatments or production processes. Later, in Section 2.3 we return to the meaning of definition (4) for normal and, more generally, LS-models.

While $\gamma(F, G)$ is a natural measure of deviation from stochastic order, it is not the only possible choice. In Leshno and Levy (2002), the authors introduce the so-called almost stochastic dominance which, easily, leads to an index to measure this deviation defining

$$\alpha(F, G) = \frac{\int_{\{G>F\}} (G(x) - F(x)) dx}{\int_{-\infty}^{\infty} |F(x) - G(x)| dx}.$$

Although $\gamma(F, G)$ is well defined for any pair of d.f.'s, some assumptions on F and G are needed in the case of $\alpha(F, G)$. In Leshno and Levy (2002), the authors require the distributions to be bounded [and limit themselves to the cases in which $\alpha(F, G) < 0.5]$. This can be relaxed, but F and G should at least have finite mean.

This index enjoys nice properties related to the expectation of utility functions (see also Tsetlin et al., 2015). However, it lacks an important property: the stochastic dominance is preserved by monotone functions. This remains true for $\gamma(F, G)$ but not for $\alpha(F, G)$. Returning again to the athletes example, let h be a strictly increasing function, and assume that we decide to measure the performance of the athletes using the values of $h(h_{old})$ and $h(h_{new})$. If F_h and G_h represent the distribution functions of the new r.v.'s, then $\gamma(F, G) = \gamma(F_h, G_h)$, while, there is no guarantee that $\alpha(F, G) = \alpha(F_h, G_h)$. We do not pursue further the analysis of the α index in this paper.

Another alternative approach to measure agreement with stochastic order has been introduced in Álvarez-Esteban et al. (2016) (see also Álvarez-Esteban et al., 2014). It is based on looking for statistical evidence supporting that, for a given (small enough) π , there exist mixture decompositions:

$$(5) \quad \begin{cases} F = (1 - \pi)\tilde{F} + \pi H_F, \\ G = (1 - \pi)\tilde{G} + \pi H_G, \end{cases} \quad \text{for some d.f.'s } \tilde{F}, \tilde{G} \text{ such that } \tilde{F} \leq_{st} \tilde{G}.$$

We mention some facts in favor of this approach. First, it allows a robust treatment of the problem of stochastic dominance because the decompositions above can be interpreted as contamination neighborhoods of some latent distributions \tilde{F} and \tilde{G} . In this sense, we should recall that statistical practitioners often process the samples to avoid “rarities” or noise; hence, a methodological approach including an adequate treatment for this kind of procedure is helpful. On the other hand, by taking π large enough, model (5) always holds (the extreme choice $\pi = 1$ will always do). The smallest π for which a such decomposition is possible measures the fraction of the population intrinsically outside the stochastic order model. This provides an index of disagreement with the stochastic order model, similar to the lack of fit index introduced in Rudas, Clogg and Lindsay (1994) for multinomial models or in Álvarez-Esteban et al. (2012) as a relaxation of the homogeneity model. The key fact to use the contamination model to measure deviation from stochastic order is given by the following result, which is contained in Álvarez-Esteban et al. (2016).

PROPOSITION 2.1. For arbitrary d.f.'s, F, G , and $\pi \in [0, 1)$, (5) holds if and only if $\pi \geq \pi(F, G)$, where

$$(6) \quad \pi(F, G) := \sup_{x \in \mathbb{R}} (G(x) - F(x)).$$

Notice that, when $\pi(F, G) \in (0, 1)$ if F and G have continuous densities f and g , respectively, then there exists x_0 such that $\pi(F, G) = G(x_0) - F(x_0)$ and x_0 satisfies $f(x_0) = g(x_0)$. Also, as $\gamma(F, G)$, $\pi(F, G)$ is invariant for strictly increasing transformations (see Remark 2.6.1 in Álvarez-Esteban et al., 2016).

A better insight into the meaning of model (5) is gained through the idea of trimmed probabilities. An

α -trimming of a probability, P , is any other probability, say Q , such that

$$Q(A) = \int_A \tau dP$$

for some function τ taking values in $[0, \frac{1}{1-\alpha}]$ and every event A . The role of the function τ is to allow to discard or downplay the influence of some regions on the sample space (having probability up to $1 - \alpha$) on the model, mimicking the common use in robust statistics of removing disturbing observations. Identifying a probability on the line with its d.f., if we write $\mathcal{R}_\alpha(F)$ for the set of trimmings of F , then $F = (1 - \alpha)\tilde{F} + \alpha H_F$ for some d.f. H_F if and only if $\tilde{F} \in \mathcal{R}_\alpha(F)$. Hence, model (5) holds if and only if there exist $\tilde{F} \in \mathcal{R}_\alpha(F)$ and $\tilde{G} \in \mathcal{R}_\alpha(G)$ such that $\tilde{F} \leq_{st} \tilde{G}$. Even more, this happens if and only if, after trimming the right tail of F and the left tail of G (removing an α fraction in both cases), the resulting \tilde{F} and \tilde{G} satisfy $\tilde{F} \leq_{st} \tilde{G}$. We refer to Álvarez-Esteban et al. (2016) for details.

Figure 4 shows the comparison of the empirical distributions corresponding to two samples of heights of boys and girls (12 years old) from a data set discussed in Álvarez-Esteban et al. (2014). The value $\pi(G_m, F_n) = 0.046$, means that it suffices to trim the fraction $\pi = 0.046$ of shortest girls and of taller boys to achieve stochastic order between the trimmed distributions and shows, up to 0.046 contamination, girls (in the sample) are taller than boys at age 12. On the other hand, to get the stochastic dominance of boys over girls we should allow a considerably higher contamination level [because $\pi(F_n, G_m) = 0.123$]; see Álvarez-Esteban et al. (2014) for a more detailed analysis of this problem.




FIG. 4. Sample d.f.'s of the heights of boys, G_m , and girls, F_n , aged 12 in the NHANES dataset.




FIG. 5. Contour-plot of $\pi(N(0, 1), N(\mu, \sigma^2))$ as in (6) for different values of μ (X-axis) and σ (Y-axis).

Coming back to the role of $\pi(F, G)$ as a measure of agreement with the stochastic order, the case of the normal model is shown in the contour-plot in Figure 5. As in the case of $\gamma(F, G)$ we see that an increase in the mean is compatible with a high level of disagreement with respect to stochastic order, that is, with an improvement under the new treatment. Thus, testing the null assumption that $\pi(F_{N(\mu_1, \sigma_1^2)}, F_{N(\mu_2, \sigma_2^2)}) \geq 0.05$ against the alternative $\pi(F_{N(\mu_1, \sigma_1^2)}, F_{N(\mu_2, \sigma_2^2)}) < 0.05$ would allow to conclude, upon rejection, that the second treatment results in improvement if we are willing to remove 5% of observations on each side, while no similar conclusion would be obtained from testing $\mu_2 \leq \mu_1$ versus $\mu_2 > \mu_1$. In Section 3, we will analyze this possibility in the light of the testing procedure developed in Álvarez-Esteban et al. (2016).

A simple comparison between the deviation indices defined in (4) and (6) arises from the following simple observation. For every couple (X, Y) of random variables with marginal d.f.'s F, G , we have

$$\begin{aligned} G(x) &= P(Y \leq x, X \leq Y) + P(Y \leq x, X > Y) \\ &\leq P(X \leq x, X \leq Y) + P(X > Y) \\ &\leq F(x) + P(X > Y), \end{aligned}$$

thus considering the quantile representations $(X, Y) = (F^{-1}, G^{-1})$, since $P(X > Y) = \ell(F^{-1} > G^{-1}) = \gamma(F, G)$, from Proposition 2.1 we get the following statement.

PROPOSITION 2.2. For any pair of d.f.'s, F and G ,

$$\pi(F, G) \leq \gamma(F, G).$$

Now, both $\pi(F, G)$ and $\gamma(F, G)$ are deviation indices from the stochastic order model, taking values in $[0, 1]$ and such that $\gamma(F, G) = 0$ if and only if $\pi(F, G) = 0$, with any of these identities being also equivalent to $F \leq_{st} G$. Later, in Section 3 below, we show how to consistently test $H_0 : \pi(F, G) \geq \delta_0$ against $H_a : \pi(F, G) < \delta_0$, which, in case of rejection, would provide statistical evidence that stochastic order holds approximately. Under some additional assumptions, we will also provide inferential methods for reaching the more restrictive conclusion (in view of Proposition 2.2) that $\gamma(F, G) < \delta_0$.

To conclude this section, we would like to mention that $\pi(F, G)$ and $\gamma(F, G)$ are related to the concepts of trimming and Winsorizing, respectively. Winsorizing and trimming are very popular robustification procedures in data analysis, designed to avoid an excessive influence of the tails, mainly in presence of outliers. Recall that $\pi(F, G) \leq \delta_0$ if and only if the d.f.'s \tilde{F} and \tilde{G} that we obtain from F and G after removing the δ_0 fraction of the upper tail of F and of the lower tail of G , respectively, satisfy $\tilde{F} \leq_{st} \tilde{G}$. Winsorizing, in turn, consists in replacing the tails of the distribution with the percentile value from each end. If we assume that $F^{-1}(t) \leq G^{-1}(t)$ for t in some interval $(\gamma_1, 1 - \gamma_2)$ then $\gamma(F, G) \leq \delta_0$ if $\gamma_1 + \gamma_2 \leq \delta_0$. Of course, $F^{-1} \leq G^{-1}$ in $(\gamma_1, 1 - \gamma_2)$ if and only if the d.f.'s \tilde{F} and \tilde{G} that we obtain from F and G by

Winsorizing (both at quantiles γ_1 and $1 - \gamma_2$) satisfy $\tilde{F} \leq_{st} \tilde{G}$.

2.3 Stochastic Order and Location-Scale Models

In this subsection, we will specialize our analysis to the comparisons under a LS-model. In fact, often, and particularly in the simulation study below, we will focus on a normal model.

Let \mathcal{F}_0 be any d.f. on the real line. A d.f. F is said to belong to the LS-model based on \mathcal{F}_0 if it satisfies $F(x) = \mathcal{F}_0(\frac{x-\theta}{\lambda}), x \in \mathbb{R}$ for some “location”, $\theta \in \mathbb{R}$, and “scale”, $\lambda > 0$ parameters. The dependence on these parameters will be included in the notation through the corresponding subindices in the way $F_{\theta,\lambda}$. By resorting to the quantile functions, we obtain the characterization

$$F_{\theta,\lambda}^{-1}(y) = \lambda \mathcal{F}_0^{-1}(y) + \theta, \quad \text{for every } y \in (0, 1),$$

from which the stochastic order $F_{\theta_1,\lambda_1} \leq_{st} F_{\theta_2,\lambda_2}$ would require

$$(7) \quad (\lambda_1 - \lambda_2) \mathcal{F}_0^{-1}(y) \leq \theta_2 - \theta_1 \quad \text{for every } y \in (0, 1).$$

Under the assumption that \mathcal{F}_0 is continuous and strictly increasing (something that we will assume from now on in the LS-model setup), condition (7) holds if and only if $\lambda_1 = \lambda_2$ and $\theta_1 \leq \theta_2$. Moreover, two quantile functions in the LS-family have a crossing point, say y_0 , if and only if $\lambda_1 \neq \lambda_2$ and $\mathcal{F}_0^{-1}(y_0) = \frac{\theta_2 - \theta_1}{\lambda_1 - \lambda_2}$, therefore, if the crossing point exists, it is unique. In other words, the set $\{y : F_{\theta_1,\lambda_1}^{-1}(y) \leq F_{\theta_2,\lambda_2}^{-1}(y)\}$ is $(0, \mathcal{F}_0(\frac{\theta_2 - \theta_1}{\lambda_1 - \lambda_2})]$ or $[\mathcal{F}_0(\frac{\theta_2 - \theta_1}{\lambda_1 - \lambda_2}), 1)$ and $\ell(\{F_{\theta_1,\lambda_1}^{-1} > F_{\theta_2,\lambda_2}^{-1}\})$ is $1 - \mathcal{F}_0(\frac{\theta_2 - \theta_1}{\lambda_1 - \lambda_2})$ or $\mathcal{F}_0(\frac{\theta_2 - \theta_1}{\lambda_1 - \lambda_2})$ (depending of the sign of $\lambda_1 - \lambda_2$). Moreover, note that

$$(8) \quad \gamma(F_{\theta_1,\lambda_1}, F_{\theta_2,\lambda_2}) = \gamma(F_{0,1}, F_{\frac{\theta_2 - \theta_1}{\lambda_1}, \frac{\lambda_2}{\lambda_1}}),$$

hence, we can focus our analysis on comparison to the reference d.f., \mathcal{F}_0 .

Now, given two d.f.’s F, G in the LS-model, if we are interested in guaranteeing an agreement with stochastic dominance of G over F of, say, 95% for the $\gamma(F, G)$ index, it would suffice to consider the crossing point of the quantile functions and check whether the interval corresponding to $\{F^{-1} \leq G^{-1}\}$ has, at least, length 0.95. For the normal model, when the reference d.f. is Φ , the standard normal d.f., a simple computation (which generalizes to any LS-model replacing Φ with the reference d.f., \mathcal{F}_0) shows that

$$\gamma(N(0, 1), N(\mu, \sigma^2)) = 1 - \Phi\left(\frac{\mu}{|\sigma - 1|}\right), \quad \sigma \neq 1,$$

while $\gamma(N(0, 1), N(\mu, 1)) = 0$ if $\mu \geq 0$ and $\gamma(N(0, 1), N(\mu, 1)) = 1$ if $\mu < 0$. This identity has been used

to obtain the contour-plot in Figure 3. We see that $\gamma(N(0, 1), N(\mu, \sigma^2))$ is constant along rays $\{(\mu, \sigma) : \mu = C|\sigma - 1|, \sigma > 0\}$, for some $C > 0$, and becomes singular at $\mu = 0, \sigma = 1$. We also note that as σ grows bigger 1 (the case of higher variance in the second sample), we can have $\mu > 0$ while $\gamma(N(0, 1), N(\mu, \sigma^2)) \rightarrow \frac{1}{2}$. This shows again that the conclusion $\mu_{new} > \mu_{old}$ is compatible with a worse performance with the new treatment for up to 50% of the population.

We analyze now the behavior of $\pi(F, G)$ under the LS-model. The fact that

$$(9) \quad \begin{aligned} & \sup_{x \in \mathbb{R}} (F_{\theta_2,\lambda_2}(x) - F_{\theta_1,\lambda_1}(x)) \\ &= \sup_{x \in \mathbb{R}} (F_{\frac{\theta_2 - \theta_1}{\lambda_1}, \frac{\lambda_2}{\lambda_1}}(x) - F_{0,1}(x)), \end{aligned}$$

shows that, as in (8), we have

$$\pi(F_{\theta_1,\lambda_1}, F_{\theta_2,\lambda_2}) = \pi(F_{0,1}, F_{\frac{\theta_2 - \theta_1}{\lambda_1}, \frac{\lambda_2}{\lambda_1}}),$$

and we can consider only the case $F = F_{0,1}$. There is no simple, general expression for $\pi(F_{0,1}, F_{\theta,\lambda})$ for every LS-model, since the maximization problem in (9) depends on \mathcal{F}_0 . In the particular case of the normal model some elementary computations show that, for $\sigma \neq 1$ and $\mu \geq 0$, $\pi(N(0, 1), N(\mu, \sigma^2)) = \Phi(\frac{\tilde{x} - \mu}{\sigma}) - \Phi(\tilde{x})$, with

$$(10) \quad \tilde{x} = \frac{\mu \pm \sqrt{\mu^2 \sigma^2 + 2\sigma^2(\sigma^2 - 1) \log \sigma}}{1 - \sigma^2},$$

where the positive sign is taken for $\sigma > 1$ and the negative for $\sigma < 1$. Also note that for the same values of μ and σ , $\pi(N(\mu, \sigma^2), N(0, 1)) = \Phi(\tilde{x}) - \Phi(\frac{\tilde{x} - \mu}{\sigma})$, with \tilde{x} the other solution in (10), a fact that allows also to get the solution for nonpositive μ . Of course, when $\sigma = 1$ and $\mu \geq 0$, $\pi(N(0, 1), N(\mu, 1)) = 0$, while $\pi(N(\mu, 1), N(0, 1))$ is attained at the only crossing point of both density functions $\tilde{x} = \mu/2$. These computations have been used to produce the contour-plot in Figure 5. We see that $\pi(N(0, 1), N(\mu, \sigma^2))$ has a smoother behavior than $\gamma(N(0, 1), N(\mu, \sigma^2))$. For a better understanding of the different roles of $\pi(N(0, 1), N(\mu, \sigma^2))$ and $\gamma(N(0, 1), N(\mu, \sigma^2))$, we include Figure 6 below. We see that $\gamma(N(0, 1), N(\mu, \sigma^2))$ equals the common value of the d.f.’s at the crossing point, while $\pi(N(0, 1), N(\mu, \sigma^2))$ equals the difference between d.f.’s at the point where the density functions have a crossing point and the $N(\mu, \sigma^2)$ d.f. is above the standard normal d.f. In the next section, we will use this characterization of $\gamma(F, G)$ in terms of crossing points (following a similar approach




FIG. 6. Distribution and density functions $N(0, 1)$ (red) and $N(1, 6)$ (green). Abscissa C (resp., D) corresponds to crossing points of the distribution (resp., density) functions. Lengths of vertical lines at C and D are $\gamma(N(0, 1), N(1, 6))$ and $\pi(N(0, 1), N(1, 6))$, respectively.

to that in Hawkins and Kochar, 1991) to design a test for the null $H_0 : \gamma(F, G) \geq \delta_0$ against the alternative $H_a : \gamma(F, G) < \delta_0$. Similarly, we will discuss how to consistently test $H_0 : \pi(F, G) \geq \delta_0$ against the alternative $H_a : \pi(F, G) < \delta_0$. These will be, according to the discussion above, feasible ways to gather statistical evidence that for approximate stochastic dominance (i.e., to conclude that, essentially, the new treatment is better than the old).

3. TESTING APPROXIMATE STOCHASTIC ORDER

In this section, we will succinctly analyze feasible test procedures to provide statistical evidence of stochastic dominance. We keep in mind that no valid inferential procedure can handle the testing problem (3) and, consequently, we have to settle with the less ambitious testing problems:

$$(11) \quad \begin{aligned} H_0 : \pi(F, G) \geq \pi_0 & \quad \text{against the alternative} \\ H_a : \pi(F, G) < \pi_0, \end{aligned}$$

or

$$(12) \quad \begin{aligned} H_0 : \gamma(F, G) \geq \gamma_0 & \quad \text{against the alternative} \\ H_a : \gamma(F, G) < \gamma_0, \end{aligned}$$

with $\pi(F, G)$ and $\gamma(F, G)$ defined in (6) and (4), respectively. We insist that rejection of the null in (11) would provide statistical evidence that stochastic order, up to some small contamination, holds. In (12), rejection would lead to guarantee, at the desired level, that treatment G produces better results than treatment F for at least a fraction of size $1 - \gamma_0$ of the population.

For this, we will first include some results, obtained in Álvarez-Esteban et al. (2016), that allow to tackle (11) and present some new results for dealing with (12) in a purely nonparametric way. Moreover, we will include some variations that can be used in the normal model. We will finish giving a comparative analysis under the normal model, based on a simulation study, providing a picture of the feasibility and performance of the different approaches. In both (11) and (12), resorting to the usual duality between one sided testing and confidence bounds, we would be interested in obtaining an upper confidence bound, say for $\pi(F, G)$, say for $\gamma(F, G)$. If $\hat{U} = \hat{U}(X_1, \dots, X_n, Y_1, \dots, Y_m)$ (resp., \hat{V}) were an (asymptotic) upper confidence bound for $\pi(F, G)$ [resp., for $\gamma(F, G)$], rejection of H_0 in (11) [resp., in (12)] when $\hat{U} < \pi_0$ (resp., when $\hat{V} < \gamma_0$) would yield a test with (asymptotically) controlled type I error probability. This will be done in two different settings corresponding to nonparametric and parametric points of view.

Throughout the section, we will assume that X_1, \dots, X_n and Y_1, \dots, Y_m are independent i.i.d. r.v.'s obtained from the d.f.'s F and G , respectively. We recall that F_n and G_m will denote the sample distribution functions based on the X 's and Y 's samples. As a common assumption in both setups, we will suppose that

$$(13) \quad \begin{aligned} & F \text{ and } G \text{ are continuous;} \\ & n, m \rightarrow \infty, \quad \lambda_{n,m} := \frac{n}{n+m} \rightarrow \lambda \in (0, 1). \end{aligned}$$

3.1 Testing Approximate Stochastic Dominance with the π Index

The role of $\pi(F, G)$ in Proposition 2.1 suggests addressing the testing problem (11) on the basis of the Kolmogorov–Smirnov statistic:

$$\pi(F_n, G_m) = \sup_{x \in \mathbb{R}} (G_m(x) - F_n(x)).$$

Consistency in the strong sense follows from the Glivenko–Cantelli theorem, which implies

$$\pi(F_n, G_m) \rightarrow \pi(F, G) \quad \text{almost surely, as } m, n \rightarrow \infty,$$

while the asymptotic behavior of its law, under assumption (13), was obtained by Raghavachari (1973), extending well-known results by Kolmogorov and Smirnov, in the way

$$(14) \quad \sqrt{\frac{mn}{m+n}} (\pi(F_n, G_m) - \pi(F, G)) \rightarrow_w \bar{B}(F, G, \lambda).$$

The limit law in (14) is that of the maximization of a Gaussian Process on a suitable set

$$\begin{aligned} \bar{B}(F, G, \lambda) := & \sup_{t \in T(F, G, \pi(F, G))} (\sqrt{\lambda} B_1(t) \\ & - \sqrt{1-\lambda} B_2(t - \pi(F, G))), \end{aligned}$$

where $B_1(t)$ and $B_2(t)$ are i.i.d. Brownian Bridges on $[0, 1]$, and the set is

$$(15) \quad T(F, G, \pi) := \{t \in [\pi, 1] : G(x) = t \text{ and } F(x) = t - \pi \text{ for some } x \in \bar{\mathbb{R}}\}.$$

Also Álvarez-Esteban et al. (2016) provides new results, involving bootstrap and computational issues, as well as exponential bounds for both types of error probabilities in testing on the basis of this statistic. Moreover (see the extended version Álvarez-Esteban et al., 2014), it is shown that, for $\alpha \in (0, 1/2)$ there are sharp lower bounds for the α -quantiles of the law of $\bar{B}(F, G, \lambda)$, given by the quantiles of a “least favorable” normal law $N(0, \bar{\sigma}_{\pi(F, G)}^2(\lambda))$ depending of λ and of $\pi(F, G)$.

In fact [in the interesting cases where $\lambda\pi \leq \frac{1}{2}$ and $(1-\lambda)\pi \leq \frac{1}{2}$], taking $\bar{\sigma}_{\pi}^2(\lambda) = \frac{1}{4} - \pi^2\lambda(1-\lambda)$, rejection of the null in (11) when

$$(16) \quad \sqrt{\frac{nm}{n+m}} (\pi(F_n, G_m) - \pi_0) < \bar{\sigma}_{\pi_0}(\lambda_{n,m}) \Phi^{-1}(\alpha),$$

provides a test of uniform asymptotic level α : If $\mathbb{P}_{F, G}$ denotes for the probability when the samples are obtained from F and G , then

$$\lim_{n \rightarrow \infty} \sup_{(F, G) \in H_0} \mathbb{P}_{F, G} \left[\sqrt{\frac{nm}{n+m}} (\pi_{n,m} - \pi_0) < \bar{\sigma}_{\pi_0}(\lambda_{n,m}) \Phi^{-1}(\alpha) \right] = \alpha.$$

Moreover, it detects alternatives with power exponentially close to one (see Proposition 3.3 in Álvarez-Esteban et al., 2016).

Two modifications of (16) result in improving the finite sample performance of the test. The first relies on the (classical) substitution of the least favorable variance by an appropriate estimation. The second, with a more important effect, tries to correct the intrinsic bias of $\pi(F_n, G_m)$, a task that often is successfully carried by resorting to the average of a set of bootstrap estimates. The final proposal, based on these modifications, $\hat{\pi}_{n,m, \text{BOOT}}$ of $\pi(F_n, G_m)$ and $\hat{\sigma}_{n,m}$ of $\bar{\sigma}_{\pi_0}(\lambda_{n,m})$ (see the details in Álvarez-Esteban et al., 2016), consists in rejection of $H_0 : \pi(F, G) \geq \pi_0$ if

$$(17) \quad \sqrt{\frac{nm}{n+m}} (\hat{\pi}_{n,m, \text{BOOT}} - \pi_0) < \hat{\sigma}_{n,m} \Phi^{-1}(\alpha),$$

which defines a test of asymptotic level α with quickly decreasing type I and type II error probabilities away from the null hypothesis boundary. Of course, by defining

$$(18) \quad \hat{U} := \hat{\pi}_{n,m, \text{BOOT}} - \sqrt{\frac{n+m}{nm}} \hat{\sigma}_{n,m} \Phi^{-1}(\alpha)$$

we get an upper bound with asymptotic confidence level at least $1 - \alpha$ for $\pi(F, G)$.

Let us notice that Section 3.4 in Álvarez-Esteban et al. (2016) is devoted to the adaptation of these statistical tools to the dependent data setup.

3.2 Testing Approximate Stochastic Dominance with the γ Index

In this subsection, we introduce a new procedure for the testing problem (12) under some additional assumption on F and G . Basically, we assume that the d.f.’s F and G have a single crossing point. This kind of assumption has been considered in some other setups. Hawkins and Kochar (1991) (see also Chen, Chen and Chen, 2002) proposed and analyzed an estimator of this crossing point, x^* . They also stress on the interest of this point for comparison of lifetimes under treatments, because, if the r.v. of interest is the survival

time, then x^* is the threshold such that, say, the control subjects have a lower chance of survival to any age $x < x^*$, while they have a higher chance of survival to ages $x > x^*$. Here, we will make the same assumption, but since our interest concerns the common value $\gamma^* = F(x^*) = G(x^*)$ at this point, we will state it in terms of the quantile functions in the alternative way: there is a unique $\gamma^* \in (0, 1)$ such that $F^{-1}(\gamma^*) = G^{-1}(\gamma^*)$ and

$$(19) \quad \begin{aligned} &F^{-1}(t) - G^{-1}(t) \quad \text{has opposite signs on} \\ &(0, \gamma^*) \quad \text{and on } (\gamma^*, 1). \end{aligned}$$

In the same spirit as in Hawkins and Kochar (1991), we introduce

$$\begin{aligned} \psi(\gamma) &= \int_0^\gamma (F^{-1}(t) - G^{-1}(t)) dt \\ &\quad - \int_\gamma^1 (F^{-1}(t) - G^{-1}(t)) dt \\ &= 2 \int_0^\gamma (F^{-1}(t) - G^{-1}(t)) dt - (\mu_F - \mu_G), \end{aligned}$$

where μ_F, μ_G denote the means of F and G . The following result gives the link between ψ and $\gamma(F, G)$.

PROPOSITION 3.1. *If F and G have finite means and positive density functions on \mathbb{R} and satisfy (19), then γ^* is the unique maximizer of $|\psi(t)|$ on $(0, 1)$, $\psi(\gamma^*) \neq 0$ and*

$$\gamma(F, G) = \begin{cases} \gamma^* & \text{if } \psi(\gamma^*) > 0, \\ 1 - \gamma^* & \text{if } \psi(\gamma^*) < 0. \end{cases}$$

This suggests that we consider to estimate γ^* by

$$\gamma_{n,m}^* = \min_{\gamma \in (0,1)} (\text{argmax}_{\gamma \in (0,1)} |\psi_{n,m}(\gamma)|),$$

with $\psi_{n,m}(\gamma) = 2 \int_0^\gamma (F_n^{-1}(t) - G_m^{-1}(t)) dt - (\bar{X}_n - \bar{Y}_m)$, and $\gamma(F, G)$ by

$$\hat{\gamma}_{n,m} = \begin{cases} \gamma_{n,m}^* & \text{if } \psi_{n,m}(\gamma_{n,m}^*) \geq 0, \\ 1 - \gamma_{n,m}^* & \text{if } \psi_{n,m}(\gamma_{n,m}^*) < 0. \end{cases}$$

The asymptotic behavior of $\hat{\gamma}_{n,m}$ is given next.

PROPOSITION 3.2. *Assume that F and G satisfy the conditions of Proposition 3.1 with densities, f and g which are continuous in a neighborhood of x^* , and such that $f(x^*) \neq g(x^*)$. Then, if the sample sizes satisfy (13),*

$$\sqrt{\frac{nm}{n+m}} (\hat{\gamma}_{n,m} - \gamma(F, G)) \rightarrow_w N(0, \sigma^2),$$

where

$$\sigma^2 = \frac{\gamma^*(1 - \gamma^*)[(1 - \lambda)g^2(x^*) + \lambda f^2(x^*)]}{(g(x^*) - f(x^*))^2}.$$

We can now base our rejection rule for (12) on a bootstrap estimation of σ^2 . More precisely, we will reject $H_0 : \gamma(F, G) \geq \gamma_0$ if

$$(20) \quad \sqrt{\frac{nm}{n+m}} (\hat{\gamma}_{n,m} - \gamma_0) < \hat{\sigma}_{n,m} \Phi^{-1}(\alpha),$$

where $\hat{\sigma}_{n,m}$ is the bootstrap estimator of σ . This rejection rule provides a consistent test of asymptotic level α . Also, as in (18),

$$(21) \quad \hat{V} := \hat{\gamma}_{n,m} - \sqrt{\frac{n+m}{nm}} \hat{\sigma}_{n,m} \Phi^{-1}(\alpha)$$

provides an upper confidence bound for $\gamma(F, G)$ with asymptotic level $1 - \alpha$.

REMARK 3.3. We must to point out the singularity of the procedure if the density functions coincide at the cross point of the d.f.'s. This fact produces instability of the approach when the densities are very similar. In particular, under a LS-model this will happen when the scales are very similar, a case that should lead to guarantee an stochastic dominance on the basis of the estimates of the location parameters.

3.3 The Parametric Point of View

If we assume that F and G belong to the same LS-model, with \mathcal{F}_0 continuous and strictly increasing, then the values $\pi(F, G)$ and $\gamma(F, G)$ can be obtained from the parameters. Therefore, the considered nonparametric approaches have natural competitors based on the estimation of the parameters. Of course, such parametric alternatives will be highly nonrobust, specially for small values of the population sizes, that are the interesting ones but strongly depend on the tails of the distributions. However, we will consider these parametric alternatives in order to explore the relative performance of our above proposals under perfect conditions. In this subsection, we assume that the parent distribution functions F, G are respectively $F_{N(\mu_1, \sigma_1^2)}, F_{N(\mu_2, \sigma_2^2)}$, although other parametric LS-families could be treated in the same way.

By considering the maximum likelihood estimators \bar{X}_n, \bar{Y}_m and S_X^2, S_Y^2 for the involved parameters, we can use the plugin estimators:

$$(22) \quad \hat{\pi}_{n,m} := \pi(F_{N(\bar{X}_n, S_X^2)}, F_{N(\bar{Y}_m, S_Y^2)}) \quad \text{for } \pi(F, G),$$

$$(23) \quad \hat{\gamma}_{n,m} := \gamma(F_{N(\bar{X}_n, S_X^2)}, F_{N(\bar{Y}_m, S_Y^2)}) \quad \text{for } \gamma(F, G).$$

Since these estimators are differentiable functions of the means and the standard deviations of the samples (excepting when $\sigma_1 = \sigma_2$, for γ), they will be asymptotically normal with the possible exception of γ in the case $\sigma_1 = \sigma_2$. Avoiding this case, the consistency and asymptotic normality would be guaranteed, and once more the bootstrap can be used to approximate the distributions of (22) and (23). The derivation of the tests and upper bounds would parallel those in Sections 3.1 and 3.2.

3.4 Some Simulations

We present now a simulation study that shows the power of the procedures discussed in Sections 3.1 and 3.2 for the assessment of approximate stochastic order. In our simulations, we have generated pairs of independent i.i.d. random samples of sizes 100, 1000 and 5000 and report the rejection frequencies of the null hypotheses (11) and (12) for three choices of π_0 and γ_0 ($\pi_0, \gamma_0 \in \{0.01, 0.05, 0.1\}$). In all cases, we have chosen the underlying distribution functions, F and G , to be normal. This allows to compare the performance of the nonparametric tests (16) (with the bootstrap bias correction) and (20) to the parametric procedures discussed in Section 3.3. Needless to say, these

parametric procedures are inconsistent if the unknown random generators F and G do not exactly fit into the parametric model. On the other hand, if F and G satisfy the parametric model then the parametric procedures should be more efficient. Hence, the performance of these parametric procedures should be taken as an ideal benchmark to which we compare the performance of the nonparametric, consistent procedures. More extensive simulations, showing the performance of the test (16) in different setups, including the least favorable cases, can be found in the Online Appendix 2 to [Álvarez-Esteban et al. \(2016\)](#).

The results are reported in Tables 1 and 2. In all cases, the nominal level of the test is $\alpha = 0.05$, F is the $N(0, 1)$ distribution and G is a $N(\mu, \sigma^2)$. Table 1 deals with the testing problem (11). Several choices of σ ($\sigma \in \{0.7, 1, 1.5\}$) have been considered. For each of these three choices of sigma, there are three different values of μ , chosen to make $\pi(F, G) = 0.01$ (left column), $\pi(F, G) = 0.05$ (central column) and $\pi(F, G) = 0.1$ (right column). Then, for each combination of sample sizes, of parameters, μ and σ , and of tolerance level, π_0 , there are two reported rejection frequencies, with the upper row corresponding to the nonparametric test and the lower row to the parametric test.

TABLE 1

Rejection rates for $\pi(N(0, 1), N(\mu, \sigma^2)) \geq \pi_0$ at the level $\alpha = 0.05$ along 1000 simulations. Upper (resp., lower) rows show the results for nonparametric (resp., parametric) comparisons. The means for each σ have been chosen to satisfy $\pi(N(0, 1), N(\mu, \sigma^2))$ equal to 0.01, 0.05 and 0.10 (resp., first, second and third columns)

π_0	Sample size	$\sigma = 0.7$ means			$\sigma = 1$ means			$\sigma = 1.5$ means		
		0.443	0.143	-0.050	-0.025	-0.125	-0.251	0.770	0.287	-0.017
0.01	100	0.132	0.010	0.000	0.026	0.011	0.001	0.110	0.009	0.000
		0.095	0.001	0.000	0.002	0.000	0.000	0.113	0.002	0.000
	1000	0.067	0.000	0.000	0.018	0.000	0.000	0.048	0.000	0.000
		0.085	0.000	0.000	0.044	0.001	0.000	0.080	0.000	0.000
	5000	0.044	0.000	0.000	0.016	0.000	0.000	0.049	0.000	0.000
		0.062	0.000	0.000	0.072	0.000	0.000	0.069	0.000	0.000
0.05	100	0.379	0.069	0.005	0.071	0.020	0.006	0.431	0.059	0.003
		0.675	0.096	0.008	0.230	0.084	0.016	0.730	0.103	0.007
	1000	0.993	0.031	0.000	0.397	0.015	0.000	0.997	0.065	0.000
		1.000	0.051	0.000	0.737	0.060	0.000	1.000	0.083	0.000
	5000	1.000	0.035	0.000	0.979	0.028	0.000	1.000	0.037	0.000
		1.000	0.057	0.000	1.000	0.057	0.000	1.000	0.052	0.000
0.10	100	0.788	0.270	0.033	0.222	0.087	0.027	0.822	0.247	0.046
		0.960	0.450	0.083	0.543	0.290	0.088	0.978	0.489	0.077
	1000	1.000	0.934	0.035	0.990	0.615	0.028	1.000	0.967	0.046
		1.000	0.992	0.059	1.000	0.877	0.053	1.000	0.996	0.061
	5000	1.000	1.000	0.029	1.000	0.998	0.032	1.000	1.000	0.039
		1.000	1.000	0.055	1.000	1.000	0.049	1.000	1.000	0.053

TABLE 2

Rejection rates for $\gamma(N(0, 1), N(\mu, \sigma^2)) \geq \gamma_0$ at the level $\alpha = 0.05$ along 1000 simulations. Upper (resp., lower) rows show the results for nonparametric (resp., parametric) comparisons. The means for each σ have been chosen to satisfy $\gamma(N(0, 1), N(\mu, \sigma^2))$ equal to 0.01, 0.05 and 0.10 (resp., first, second and third columns)

γ_0	Sample size	$\sigma = 1.1$ means			$\sigma = 1.5$ means			$\sigma = 2$ means		
		0.233	0.164	0.128	1.163	0.822	0.641	2.326	1.645	1.282
0.01	100	0.000	0.000	0.000	0.007	0.000	0.000	0.000	0.006	0.000
		0.001	0.000	0.000	0.070	0.006	0.000	0.124	0.001	0.000
	1000	0.013	0.000	0.000	0.095	0.000	0.000	0.102	0.000	0.000
		0.038	0.004	0.000	0.085	0.000	0.000	0.083	0.000	0.000
	5000	0.046	0.001	0.001	0.097	0.000	0.000	0.065	0.000	0.000
		0.097	0.003	0.000	0.076	0.000	0.000	0.048	0.000	0.000
0.05	100	0.015	0.003	0.000	0.338	0.058	0.015	0.611	0.089	0.017
		0.038	0.008	0.001	0.491	0.112	0.035	0.764	0.116	0.016
	1000	0.209	0.043	0.007	0.919	0.088	0.002	0.998	0.082	0.000
		0.319	0.079	0.015	0.992	0.089	0.001	1.000	0.066	0.000
	5000	0.654	0.084	0.008	1.000	0.048	0.000	1.000	0.044	0.000
		0.799	0.105	0.011	1.000	0.053	0.000	1.000	0.053	0.000
0.10	100	0.061	0.027	0.007	0.702	0.256	0.095	0.926	0.390	0.121
		0.090	0.035	0.016	0.810	0.310	0.141	0.978	0.461	0.112
	1000	0.540	0.212	0.073	1.000	0.661	0.084	1.000	0.884	0.066
		0.672	0.258	0.092	1.000	0.795	0.074	1.000	0.988	0.057
	5000	0.964	0.395	0.105	1.000	0.993	0.060	1.000	1.000	0.049
		0.987	0.437	0.102	1.000	0.999	0.062	1.000	1.000	0.055

We see that the rejection frequencies for the nonparametric procedure show either a reasonable agreement to the nominal level of the test or are slightly conservative, while the parametric procedure is slightly liberal. On the other hand, the nonparametric procedure is able to reject the null with remarkably high power. As an illustration, consider, for instance, the block $\sigma = 0.7, \pi_0 = 0.05$. Within this block, the boundary between the null and the alternative hypotheses corresponds to the middle column [$\mu = 0.143$; then $\pi(F, G) = 0.05$]. The observed rejection frequencies for the nonparametric procedure are 0.031 and 0.035 for sample sizes $n = m = 1000$ and $n = m = 5000$, respectively, a bit below the nominal level of the test ($\alpha = 0.05$). As we move to the alternative [the left column, $\mu = 0.443, \pi(F, G) = 0.01$], we see that samples of size $n = m = 1000$ are enough to reject the null hypothesis $\pi(F, G) \geq 0.05$ (and conclude that F and G satisfy stochastic order up to less than 5% contamination) with high probability (the observed rejection frequency is 0.993). The worst behaviour in terms of power corresponds to the case $\sigma = 1$ (middle group). In this case, rejection of the null with high power (90% or higher) requires sample sizes $n = m = 5000$. We note,

nevertheless, that testing for approximate stochastic order is a hard inferential problem and, on the other hand, sample sizes in this range are not unusual in many fields of application. As for the parametric procedure introduced for comparison (bottom rows), we observe that it presents a better performance in terms of power but it is a bit liberal in some cases (and recall, again, that it is not a consistent procedure as we move away from this LS setup).

The results for the testing problem (12) are reported in Table 2. In this setup, the cases $\sigma = \sigma_0$ and $\sigma = 1 - \sigma_0$ are symmetrical and we have focused on the case $\sigma \geq 1$. The case $\sigma = 1$ would need a different handling, since $\gamma(N(0, 1), N(\mu, 1))$ only takes the values 0 and 1 depending on when $\mu \geq 0$ or $\mu < 0$. For these reasons, we have fixed $\sigma \in \{1.1, 1.5, 2\}$, choosing then μ accordingly to get $\gamma(N(0, 1), N(\mu, \sigma^2)) \in \{0.01, 0.05, 0.10\}$.

We see in this case that the tests (both the nonparametric and parametric) can be too liberal if the two samples have similar variances. This is not surprising, since the asymptotic variance in Proposition 3.2 tends to ∞ as $\sigma \rightarrow 1$. As we move away from this singular case, we see a somewhat better degree of agreement

to the nominal level of the test, as well as a generally good performance in terms of power.

Deviations from the ideal model of stochastic order in terms of the $\gamma(F, G)$ -index admit, arguably, a simpler interpretation than deviations in $\pi(F, G)$ -index, but we see that the assessment of stochastic order up to a small deviation in $\pi(F, G)$ -index is, from the point of view of statistical inference, a better posed problem, less affected by the similarity of variances. Finally, we remark that, although both indices are intrinsically nonparametric in nature, the scope of $\pi(F, G)$ is considerably larger, since the single crossing point assumption required for the validity of the asymptotic theory for the $\gamma(F, G)$ -index could be too restrictive for some real applications.

APPENDIX: PROOFS

To the best of our knowledge, the index $\gamma(F, G)$ has been introduced just here, thus we include in this Appendix some technical details to justify our claims about the asymptotics for the proposed estimator $\hat{\gamma}_{n,m}$.

PROOF OF PROPOSITION 3.1. We begin noting that $\psi(1) = \mu_F - \mu_G = -\psi(0)$ and also that ψ is differentiable in $(0, 1)$, with derivative $\psi'(t) = 2(F^{-1}(t) - G^{-1}(t))$.

Consider first the case

$$F^{-1}(t) > G^{-1}(t), \quad \text{for } t \in (0, \gamma^*), \text{ while}$$

$$F^{-1}(t) < G^{-1}(t), \quad \text{for } t \in (\gamma^*, 1),$$

we have $\gamma(F, G) = \gamma^*$ with $\psi'(t) > 0, t \in (0, \gamma^*)$, while $\psi'(t) < 0, t \in (\gamma^*, 1)$. In particular, γ^* is the unique maximizer of ψ .

Now, if $\psi(0) \geq 0$ then $\psi(1) \leq 0, \psi(\gamma^*) > \psi(0) \geq 0$ and $\psi(t) \in [\psi(1), \psi(\gamma^*)]$ for all $t \in (0, 1)$ and we see that γ^* is the unique maximizer of $|\psi(t)|$. If, on the other hand, $\psi(0) < 0$ then $\psi(\gamma^*) > \psi(1) > 0, \psi(t) \in [\psi(0), \psi(\gamma^*)]$ for all $t \in (0, 1)$ and, again, γ^* is the unique maximizer of $|\psi(t)|$.

In the other case $F^{-1}(t) < G^{-1}(t), t \in (0, \gamma^*)$, and $F^{-1}(t) > G^{-1}(t), t \in (\gamma^*, 1)$, we have $\gamma(F, G) = 1 - \gamma^*$ and, arguing as above, we see that γ^* is the unique minimizer of ψ and the unique maximizer of $|\psi(t)|$ and satisfies $\psi(\gamma^*) < 0$. \square

By considering the quantile functions F_n^{-1} associated to the empirical d.f.'s F_n as a random function, we get in a natural way the quantile process defined by $u_n^F(t) = \sqrt{n}(F_n^{-1}(t) - F^{-1}(t))$ for $t \in (0, 1)$. The study of this statistically meaningful stochastic process was addressed in the second half of the past century.

For use in the proof of Proposition 3.2, we provide the following lemma.

LEMMA A.1. *Let F (resp., G) be a d.f. with continuous and positive derivative f (resp., g) on the interval $[F^{-1}(p) - \varepsilon, F^{-1}(q) + \varepsilon]$ (resp., $[G^{-1}(p) - \varepsilon, G^{-1}(q) + \varepsilon]$). Under the independence assumption on the samples obtained from F and G , let $u_n^F(\cdot)$ and $u_m^G(\cdot)$ be the corresponding quantile processes. Then there exist independent versions $\tilde{u}_n^F(\cdot)$ and $\tilde{u}_m^G(\cdot)$, of these processes (with the same joint distribution that the originals) and independent standard Brownian bridges \tilde{B}_1 and \tilde{B}_2 , such that*

$$\sup_{t \in [p, q]} \left| \left(\sqrt{\frac{m}{n+m}} \tilde{u}_n^F(t) - \sqrt{\frac{n}{n+m}} \tilde{u}_m^G(t) \right) - \left(\sqrt{1-\lambda} \frac{B_1(t)}{f(F^{-1}(t))} - \sqrt{\lambda} \frac{B_2(t)}{g(G^{-1}(t))} \right) \right| \xrightarrow{a.s.} 0.$$

PROOF. The hypotheses on F and G guarantee (see Example 3.9.24 in van der Vaart and Wellner, 1996):

$$u_n^F(\cdot) \xrightarrow{w} \frac{B_1(\cdot)}{f(F^{-1}(\cdot))} \quad \text{and}$$

$$u_m^G(\cdot) \xrightarrow{w} \frac{B_2(\cdot)}{g(G^{-1}(\cdot))} \quad \text{in the space } L^\infty[p, q],$$

where B_1, B_2 are independent standard Brownian bridges. Moreover, the independence of the samples implies that of the quantile processes, thus the joint convergence

$$(u_n^F(\cdot), u_m^G(\cdot)) \xrightarrow{w} \left(\frac{B_1(\cdot)}{f(F^{-1}(\cdot))}, \frac{B_2(\cdot)}{g(G^{-1}(\cdot))} \right).$$

Now we can resort to the Skorohod–Dudley–Wichura almost surely representation theorem (see, e.g., Theorem 1.10.3 in van der Vaart and Wellner, 1996), providing a sequence of pairs $(\tilde{u}_n^F(\cdot), \tilde{u}_m^G(\cdot)) \stackrel{d}{=} (u_n^F(\cdot), u_m^G(\cdot))$ and a pair $(\frac{\tilde{B}_1(\cdot)}{f(F^{-1}(\cdot))}, \frac{\tilde{B}_2(\cdot)}{g(G^{-1}(\cdot))}) \stackrel{d}{=} (\frac{B_1(\cdot)}{f(F^{-1}(\cdot))}, \frac{B_2(\cdot)}{g(G^{-1}(\cdot))})$ such that $(\tilde{u}_n^F(\cdot), \tilde{u}_m^G(\cdot)) \xrightarrow{a.s.} (\frac{\tilde{B}_1(\cdot)}{f(F^{-1}(\cdot))}, \frac{\tilde{B}_2(\cdot)}{g(G^{-1}(\cdot))})$ in the space $L^\infty[p, q]$. From here, the result is straightforward. \square

PROOF OF PROPOSITION 3.2. We note first that

$$\sup_{\gamma \in [0, 1]} |\psi_{n,m}(\gamma) - \psi(\gamma)|$$

$$\leq \int_0^1 |F_n^{-1}(t) - F^{-1}(t)| dt$$

$$+ \int_0^1 |G_m^{-1}(t) - G^{-1}(t)| dt \rightarrow 0 \quad \text{a.s.}$$

[to check that $\int_0^1 |F_n^{-1}(t) - F^{-1}(t)| dt$ vanishes asymptotically we can use the fact that it equals the Wasserstein distance between F_n and F , see [del Barrio, Giné and Matrán, 1999](#)]. As a consequence, we have that $\gamma_{m,n}^* \rightarrow \gamma^*$ a.s.

Therefore, in the case $F^{-1} - G^{-1} > 0$ in $(0, \gamma^*)$, $F^{-1} - G^{-1} < 0$ in $(\gamma^*, 1)$, we will have that a.s. $\psi_{m,n}(\gamma_{m,n}^*) > 0$ eventually and $\hat{\gamma}_{m,n} = \gamma_{m,n}^*$. Similarly, in the case $F^{-1} - G^{-1} < 0$ in $(0, \gamma^*)$, $F^{-1} - G^{-1} > 0$ in $(\gamma^*, 1)$, we will have a.s. that eventually $\hat{\gamma}_{m,n} = 1 - \gamma_{m,n}^*$. Hence, in a probability one set, eventually

$$\sqrt{\frac{n+m}{nm}}(\hat{\gamma}_{m,n} - \gamma(F, G)) = \pm \sqrt{\frac{n+m}{nm}}(\gamma_{m,n}^* - \gamma^*),$$

with the positive sign in the former case and the negative in the latter.

By symmetry of the centered normal laws, it suffices to prove the convergence for $\sqrt{\frac{n+m}{nm}}(\gamma_{m,n}^* - \gamma^*)$. From this point, we assume that we are in the case $F^{-1} - G^{-1} > 0$ in $(0, \gamma^*)$, $F^{-1} - G^{-1} < 0$ in $(\gamma^*, 1)$ and note that in this case γ^* is also the maximizer of ψ .

We note also that we can replace ψ by $\phi(\gamma) = \int_\delta^\gamma (F^{-1}(t) - G^{-1}(t)) dt$ for some fixed $\delta \in (0, 1)$ and still have that γ^* is the maximizer of $\phi(\gamma)$, $\gamma \in (\delta, \delta')$ for some other $\delta' \in (0, 1)$. We similarly set $\phi_{n,m}(\gamma) = \int_\delta^\gamma (F_n^{-1}(t) - G_m^{-1}(t)) dt$. The assumptions ensure that we can choose δ and δ' such that f and g are continuous and bounded away from 0 in (δ, δ') . Then, with the notation introduced for the quantile processes associated to the samples obtained from F and G ,

$$\begin{aligned} & \sqrt{\frac{nm}{n+m}}(\phi_{n,m}(\gamma) - \phi(\gamma)) \\ &= \int_\delta^\gamma \left(\sqrt{\frac{m}{n+m}} u_n^F(t) - \sqrt{\frac{n}{n+m}} u_m^G(t) \right) dt. \end{aligned}$$

The application of Lemma A.1 to the quantile processes in $L^\infty[\delta, \delta']$, implies that there are versions of u_n^F, u_m^G, B_1 and B_2 such that

$$\begin{aligned} & \sup_{t \in [\delta, \delta']} \left| \left(\sqrt{\frac{m}{n+m}} u_n^F(t) - \sqrt{\frac{n}{n+m}} u_m^G(t) \right) \right. \\ & \left. - \left(\sqrt{1-\lambda} \frac{B_1(t)}{f(F^{-1}(t))} - \sqrt{\lambda} \frac{B_2(t)}{g(G^{-1}(t))} \right) \right| \xrightarrow{\text{a.s.}} 0. \end{aligned}$$

From this, we conclude that for any sequence verifying $\gamma_{n,m} = \gamma^* + o_P(1)$,

$$\begin{aligned} & \sqrt{\frac{nm}{n+m}}(\phi_{n,m}(\gamma_{n,m}) - \phi(\gamma_{m,n})) \\ & - \sqrt{\frac{nm}{n+m}}(\phi_{n,m}(\gamma^*) - \phi(\gamma^*)) \end{aligned}$$

$$\begin{aligned} (24) \quad &= \int_{\gamma^*}^{\gamma_{m,n}} \left(\sqrt{\frac{m}{n+m}} u_n^F(t) - \sqrt{\frac{n}{n+m}} u_m^G(t) \right) dt \\ &= (\gamma_{m,n} - \gamma^*)Z + (\gamma_{m,n} - \gamma^*)o_P(1), \end{aligned}$$

with $Z = \sqrt{1-\lambda}B_1(\gamma^*)/f(x^*) - \sqrt{\lambda}B_2(\gamma^*)/g(x^*)$. Since we already obtained the consistency $\gamma_{m,n}^* \rightarrow \gamma^*$ a.s., we can apply now Theorem A.2 below to conclude that

$$\sqrt{\frac{nm}{n+m}}(\gamma_{n,m}^* - \gamma^*) = -\frac{Z}{\frac{1}{f(x^*)} - \frac{1}{g(x^*)}} + o_P(1),$$

which completes the proof. \square

The following theorem is a suitable version of Theorem 3.2.16 in [van der Vaart and Wellner, 1996](#) (see the final comments there leading to this simplified statement). It allows to obtain the asymptotic law of an estimator, like that involved in Proposition 3.2, based on an “argmax” procedure. This kind of argument is one of the best known tools to address the asymptotics of M-estimators (see Section 3.2.4 in [van der Vaart and Wellner, 1996](#))

THEOREM A.2 (see Theorem 3.2.16 in [van der Vaart and Wellner, 1996](#)). *Let M_n be stochastic processes indexed by an open interval $\Theta \subset \mathbb{R}$ and $M : \Theta \rightarrow \mathbb{R}$ a deterministic function. Assume that $\theta \rightarrow M(\theta)$ is twice continuously derivable at a point of maximum θ_0 with second-derivative $M''(\theta_0) \neq 0$. Suppose that for some sequence $r_n \rightarrow \infty$:*

$$\begin{aligned} & r_n(M_n(\tilde{\theta}_n) - M(\tilde{\theta}_n)) - r_n(M_n(\theta_0) - M(\theta_0)) \\ &= (\tilde{\theta}_n - \theta_0)Z + o_P(|\tilde{\theta}_n - \theta_0|), \end{aligned}$$

for every random sequence $\tilde{\theta}_n = \theta_0 + o_P(1)$ and a random variable Z . If the sequence $\hat{\theta}_n \xrightarrow{P} \theta_0$ and satisfies $M_n(\hat{\theta}_n) \geq \sup_\theta M_n(\theta) - o_P(r_n^{-2})$ for every n , then

$$r_n(\hat{\theta}_n - \theta_0) = -\frac{Z}{M''(\theta_0)} + o_P(1).$$

ACKNOWLEDGMENTS

Research partially supported by the Spanish Ministerio de Economía y Competitividad y fondos FEDER, grants MTM2014-56235-C2-1-P and MTM2014-56235-C2-2, and by Consejería de Educación de la Junta de Castilla y León, grant VA212U13.

REFERENCES

- ÁLVAREZ-ESTEBAN, P. C., DEL BARRIO, E., CUESTA-ALBERTOS, J. A. and MATRÁN, C. (2008). Trimmed comparison of distributions. *J. Amer. Statist. Assoc.* **103** 697–704. [MR2435470](#)
- ÁLVAREZ-ESTEBAN, P. C., DEL BARRIO, E., CUESTA-ALBERTOS, J. A. and MATRÁN, C. (2012). Similarity of samples and trimming. *Bernoulli* **18** 606–634. [MR2922463](#)
- ÁLVAREZ-ESTEBAN, P. C., DEL BARRIO, E., CUESTA-ALBERTOS, J. A. and MATRÁN, C. (2014). A contamination model for approximate stochastic order: Extended version. Available at <http://arxiv.org/abs/1412.1920>.
- ÁLVAREZ-ESTEBAN, P. C., DEL BARRIO, E., CUESTA-ALBERTOS, J. A. and MATRÁN, C. (2016). A contamination model for the stochastic order. *TEST* **25** 751–774. [MR3554414](#)
- ANDERSON, G. (1996). Nonparametric tests for stochastic dominance. *Econometrica* **64** 1183–1193.
- BARRETT, G. F. and DONALD, S. G. (2003). Consistent tests for stochastic dominance. *Econometrica* **71** 71–104. [MR1956856](#)
- BERGER, R. L. (1988). A nonparametric, intersection-union test for stochastic order. In *Statistical Decision Theory and Related Topics, IV, Vol. 2 (West Lafayette, Ind., 1986)* 253–264. Springer, New York. [MR0927137](#)
- CHEN, G., CHEN, J. and CHEN, Y. (2002). Statistical inference on comparing two distribution functions with a possible crossing point. *Statist. Probab. Lett.* **60** 329–341. [MR1942937](#)
- DAVIDSON, R. and DUCLOS, J.-Y. (2000). Statistical inference for stochastic dominance and for the measurement of poverty and inequality. *Econometrica* **68** 1435–1464. [MR1793365](#)
- DAVIDSON, R. and DUCLOS, J.-Y. (2013). Testing for restricted stochastic dominance. *Econometric Rev.* **32** 84–125. [MR2988921](#)
- DEL BARRIO, E., GINÉ, E. and MATRÁN, C. (1999). Central limit theorems for the Wasserstein distance between the empirical and the true distributions. *Ann. Probab.* **27** 1009–1071. [MR1698999](#)
- DETTE, H. and MUNK, A. (1998). Validation of linear regression models. *Ann. Statist.* **26** 778–800. [MR1626016](#)
- HAWKINS, D. L. and KOCHAR, S. C. (1991). Inference for the crossing point of two continuous CDFs. *Ann. Statist.* **19** 1626–1638. [MR1126342](#)
- HODGES, J. L. JR. and LEHMANN, E. L. (1954). Testing the approximate validity of statistical hypotheses. *J. Roy. Statist. Soc. Ser. B* **16** 261–268. [MR0069461](#)
- LEHMANN, E. L. (1955). Ordered families of distributions. *Ann. Math. Stat.* **26** 399–419. [MR0071684](#)
- LEHMANN, E. L. and ROJO, J. (1992). Invariant directional orderings. *Ann. Statist.* **20** 2100–2110. [MR1193328](#)
- LESHNO, M. and LEVY, H. (2002). Preferred by “All” and preferred by “Most” decision makers: Almost stochastic dominance. *Manage. Sci.* **48** 1074–1085.
- LINTON, O., MAASOUMI, E. and WHANG, Y.-J. (2005). Consistent testing for stochastic dominance under general sampling schemes. *Rev. Econ. Stud.* **72** 735–765. [MR2148141](#)
- LINTON, O., SONG, K. and WHANG, Y.-J. (2010). An improved bootstrap test of stochastic dominance. *J. Econometrics* **154** 186–202. [MR2558960](#)
- LIU, J. and LINDSAY, B. G. (2009). Building and using semi-parametric tolerance regions for parametric multinomial models. *Ann. Statist.* **37** 3644–3659. [MR2549573](#)
- MANN, H. B. and WHITNEY, D. R. (1947). On a test of whether one of two random variables is stochastically larger than the other. *Ann. Math. Stat.* **18** 50–60. [MR0022058](#)
- MCFADDEN, D. (1989). Testing for stochastic dominance. In *Studies in the Economics of Uncertainty: In Honor of Josef Hadar* (T. B. Fomby and T. K. Seo, eds.) Springer, New York.
- MÜLLER, A. and STOYAN, D. (2002). *Comparison Methods for Stochastic Models and Risks*. Wiley, Chichester. [MR1889865](#)
- MUNK, A. and CZADO, C. (1998). Nonparametric validation of similar distributions and assessment of goodness of fit. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **60** 223–241. [MR1625620](#)
- RAGHAVACHARI, M. (1973). Limiting distributions of Kolmogorov–Smirnov type statistics under the alternative. *Ann. Statist.* **1** 67–73. [MR0346976](#)
- RUDAS, T., CLOGG, C. C. and LINDSAY, B. G. (1994). A new index of fit based on mixture methods for the analysis of contingency tables. *J. Roy. Statist. Soc. Ser. B* **56** 623–639. [MR1293237](#)
- SHAKED, M. and SHANTHIKUMAR, J. G. (2007). *Stochastic Orders*. Springer, New York. [MR2265633](#)
- TSETLIN, I., WINKLER, R. L., HUANG, R. J. and TZENG, L. Y. (2015). Generalized almost stochastic dominance. *Oper. Res.* **63** 363–377. [MR3338587](#)
- VAN DER VAART, A. W. and WELLNER, J. A. (1996). *Weak Convergence and Empirical Processes*. Springer, New York. [MR1385671](#)