



Loss aversion and the welfare ranking of policy interventions[☆]

Sergio Firpo^{a,*}, Antonio F. Galvao^b, Martyna Kobus^c, Thomas Parker^d,
Pedro Rosa-Dias^e

^a Insper, Sao Paulo, Brazil

^b Department of Economics, Michigan State University, East Lansing, USA

^c Institute of Economics, Polish Academy of Sciences, Warsaw, Poland

^d Department of Economics, University of Waterloo, Waterloo, Canada

^e Department of Economics and Public Policy and CHEPI, Imperial College Business School, Imperial College London, UK

ARTICLE INFO

JEL classification:

C12

C14

I30

Keywords:

Welfare

Loss aversion

Policy evaluation

Stochastic ordering

Directional differentiability

ABSTRACT

This paper develops theoretical criteria and econometric methods to rank policy interventions in terms of welfare when individuals are loss-averse. Our new criterion for “loss aversion-sensitive dominance” defines a weak partial ordering of the distributions of policy-induced gains and losses. It applies to the class of welfare functions which model individual preferences with non-decreasing and loss-averse attitudes towards changes in outcomes. We also develop new semiparametric statistical methods to test loss aversion-sensitive dominance in practice, using nonparametric plug-in estimates; these allow inference to be conducted through a special resampling procedure. Since point-identification of the distribution of policy-induced gains and losses may require strong assumptions, we extend our comparison criteria, test statistics, and resampling procedures to the partially-identified case. We illustrate our methods with a simple empirical application to the welfare comparison of alternative income support programs in the US.

We suffer more, ... when we fall from a better to a worse situation, than we ever enjoy when we rise from a worse to a better.

[Adam Smith, The Theory of Moral Sentiments]

1. Introduction

Policy interventions often generate heterogeneous effects, giving rise to gains and losses to different individuals and sectors of society. Classically, the welfare ranking of policy interventions, conducted under the Rawlsian principle of the “veil of ignorance”, has deemed such gains and losses irrelevant: all policies that produce the same marginal distribution of outcomes should be

[☆] The authors are grateful to Pedro Carneiro, Dan Hamermesh, Hide Ichimura, Radosław Kurek, Essie Maasoumi, Piotr Miłoś, Magne Mogstad, Jim Powell, João Santos Silva, Tiemen Woutersen, and seminar participants at the University of Chicago, University of Wisconsin-Madison, USC, MSU, University of Montreal, University of California Berkeley, University of Arizona, 28th annual meeting of the Midwest Econometrics Group, 35th Meeting of the Canadian Econometric Study Group, and the 3rd edition of the Rio-Sao Paulo Econometrics Workshop for useful comments and discussions regarding this paper. Martyna Kobus acknowledges the support of the National Science Centre, Poland, the grant no. 2020/39/B/HS4/03131. Marta Schoch provided excellent research assistance. Computer programs to replicate the numerical analyses are available from the authors. This work was made possible by the facilities of the Shared Hierarchical Academic Research Computing Network (SHARCNET: <https://www.sharcnet.ca>) and Compute/Calcul Canada. All the remaining errors are ours.

* Corresponding author.

E-mail addresses: firpo@insper.edu.br (S. Firpo), agalvao@msu.edu (A.F. Galvao), mkobus@inepan.waw.pl (M. Kobus), tmpparker@uwaterloo.ca (T. Parker), p.rosa-dias@imperial.ac.uk (P. Rosa-Dias).

<https://doi.org/10.1016/j.jeconom.2023.105643>

Received 18 June 2022; Received in revised form 19 May 2023; Accepted 3 October 2023

considered equivalent for the purpose of welfare analysis (Atkinson, 1970; Roemer, 1998; Sen, 2000). However, more recent approaches have focused precisely on how different individuals are affected by a given policy (Heckman and Smith, 1998; Carneiro et al., 2001). Our paper relates to this latter approach and focuses on an important characteristic of the individual valuation of policy-induced effects: loss aversion.¹

There are two main reasons why loss aversion can be important for the welfare ranking of policy interventions. First, as shown in Carneiro et al. (2001), the individual gains and losses caused by a policy have important political economy consequences. Public support for that policy, and for the authorities that implement it, depends on the balance of gains and losses experienced and valued by different individuals in the electorate. In this context, there is mounting empirical evidence indicating that the electorate often exhibits loss aversion. This aversion to losses among constituents, in turn, drives the actions of policy makers, as documented in situations as diverse as government support to the steel industry in US trade policy, and President Trump's attempted repeal of the Affordable Care Act (Freund and Özden, 2008; Alesina and Passarelli, 2019).

Second, political economy aside, there are important situations where policy makers have strong normative reasons for incorporating loss-aversion in their own welfare ranking of public policies. This has been proposed in a range of fields. In a recent example, Eyal (2020) shows that the Hippocratic principle of "first, do no harm" has led US and EU policy makers to delay the development of Covid-19 vaccines by rejecting human challenge trials, that involve purposefully infecting a small number of vaccine trial volunteers. In this situation, policy makers placed more weight on the potential harm to a small number of volunteers, than on the vast potential benefits of accelerating the availability of a vaccine to large swathes of the population. Similarly, in the context of minimum wage legislation, Mankiw (2014) proposes that policy makers should be loss-averse in their approach to policy evaluation and adopt a "first, do no harm" principle: "*As I see it, the minimum wage and the Affordable Care Act are cases in point. Noble as they are in aspiration, they fail the do-no harm test. An increase in the minimum wage would disrupt some deals that workers and employers have made voluntarily*". Along the same lines, in the context of development economics, scholars such as Easterly (2009) have proposed a "first, do no harm" approach to foreign aid interventions in developing countries.

In this paper, we extend the toolkit available for the evaluation of policy interventions in contexts where it is sensible to incorporate loss-aversion in the welfare ranking of policy interventions. We are not proposing that *all* policies should be ranked under loss aversion sensitive criteria. Instead, we provide a methodology for conducting such a ranking in cases where aversion to losses is likely to be important. As discussed above, this may be the case either because policy makers are genuinely loss-averse, or because they know that the individuals exposed to the policy are so, and this leads them to incorporate loss aversion in policy ranking. To this end, our paper develops new testable criteria and econometric methods to rank distributions of individual policy effects, from a welfare standpoint, incorporating loss aversion. We make two main contributions to the literature.

Our first contribution is to propose loss aversion-sensitive criteria for the welfare ranking of policies. We adopt the standard welfare function approach (Atkinson, 1970): alternative policies are compared based on a welfare ranking, where social welfare is an additively separable and symmetric function of individuals' outcomes. It is well established that, for non-decreasing utility functions, this is equivalent to first-order stochastic dominance (FOSD) over distributions of policy outcomes. Analogously, our ranking is based on social *value* functions, which are additively separable and symmetric functions of individual *gains and losses*. We show that the social value function ranking with non-decreasing and loss-averse value functions (Tversky and Kahneman, 1991) is equivalent to a new concept we call loss aversion-sensitive dominance (LASD) over distributions of policy-induced gains and losses. FOSD requires that the cumulative distribution function of the dominated distribution lies everywhere above the cumulative distribution of the dominant distribution. In contrast, under LASD, the dominated cumulative distribution function must lie sufficiently above the dominant distribution function for *losses* such that the probability of potential losses cannot be compensated by a higher probability for potential gains. This is a consequence of loss-aversion. Except for the special case of a *status quo* policy (i.e. a policy of no change) where FOSD and LASD coincide, generally, as we show, LASD can be used to compare policies that are indistinguishable for FOSD.²

The LASD criterion relies on gains and losses, which under standard identification conditions can be considered *treatment effects*. It is well known that the point identification of the distribution of treatment effects may require implausible theoretical restrictions such as rank invariance of potential outcomes (Heckman et al., 1997). We thus extend our LASD criteria to a partially-identified setting and establish a sufficient condition to rank alternative policies under partial identification of the distributions of their effects. We use Makarov bounds (Makarov, 1982; Rüschendorf, 1982; Frank et al., 1987) to bound the distribution of treatment effects when the joint pre and post-policy outcome distribution is unknown. This provides a testable criterion that can be used in practice, since the marginal distribution functions from samples observed under various treatments can usually be identified and Makarov bounds only rely on marginal information for their identification.

Our second contribution is to develop statistical inference procedures to practically test the loss aversion-sensitive dominance condition using sample data. We develop statistical tests for both point-identified and partially-identified distributions of outcomes. The test procedures are designed to assess, uniformly over the two outcome distributions, whether one treatment dominates another

¹ Loss aversion is a well established empirical regularity, documented in a wide variety of contexts (Kahneman and Tversky, 1979; Samuelson and Zeckhauser, 1988; Tversky and Kahneman, 1991; Rabin and Thaler, 2001; Rick, 2011).

² The literature on stochastic dominance is vast and spans economics and mathematics - we refer the reader to, e.g., Shaked and Shanthikumar (1994) and Levy (2016) for a review. When dominance curves cross, higher order or inverse stochastic dominance criteria have been proposed. The former involves conditions on higher (typically third and fourth) order derivatives of utility function (e.g. Fishburn, 1980; Chew, 1983 to which Eeckhoudt and Schlesinger, 2006 provided interesting interpretation, whereas the latter is related to the rank-dependent theory originally proposed by Weymark, 1981 and Yaari, 1987, 1988, where social welfare functions are weighted averages of ordered outcomes with weights decreasing with the rank of the outcome (see Aaberge et al., 2021 for a recent refinement of this theory)).

in terms of the LASD criterion. Specifically, we suggest Kolmogorov–Smirnov and Cramér–von Mises test statistics that are applied to nonparametric plug-in estimates of the LASD criterion mentioned above. Inference for these statistics uses specially tailored resampling procedures. We show that our procedures control the size of tests for all probability distributions that satisfy the null hypothesis. Our tests are related to the literature on inference for stochastic dominance represented by, e.g., [Linton et al. \(2010, 2005\)](#), [Barrett and Donald \(2003\)](#) and references cited therein. [Linton et al. \(2005\)](#) is an important contribution because in addition to developing tests for stochastic dominance of arbitrary order, they propose a Prospect Theory stochastic dominance test. Their test is intended for inferring dominance among a different family of value functions than ours, namely, the focus is on risk loving for gains and risk aversion for losses (i.e. so called S-shapedness, [Kahneman and Tversky, 1979](#)), but not on loss aversion. We contribute to the literature by developing tests for loss averse-sensitive dominance, which are an alternative to standard stochastic dominance tests. Our tests widen the variety of comparisons available to empirical researchers to other criteria that encode important qualitative features of agent preferences.

The LASD criterion results in a functional inequality that depends on marginal distribution functions, and we adapt existing techniques from the literature on testing functional inequalities to test for LASD. However, in comparison with stochastic dominance tests, verifying LASD with sample data presents technical challenges for both the point- and partially-identified cases. The criterion that implies LASD of one distribution over another is more complex than the standard FOSD criterion, and hence requires a significant extension of existing procedures to justify the use of inference about loss averse-sensitive dominance with a nonparametric plug-in estimator of the LASD criterion.

In particular, the problem with using existing stochastic dominance techniques is that the mapping of distribution functions to a testable criterion is nonlinear and ill-behaved. A dominance test inherently requires uniform comparisons be made, and tractable analysis of its distribution demands regularity, in the form of differentiability, of the map between the space of distribution functions and the space of criterion functions. However, the map from pairs of distribution functions to the LASD criterion function is not differentiable. Despite this complication, we show that supremum- or L_2 -norm statistics applied to this function are just regular enough that, with some care, resampling can be used to conduct inference.

For practical implementation, we propose an inference procedure that combines standard resampling with an estimate of the way that test statistics depend on underlying data distributions, building on recent results from [Fang and Santos \(2019\)](#). We contribute to the literature on directionally differentiable test statistics with a new test for LASD. Recent contributions to this literature include, among others, [Hong and Li \(2018\)](#), [Chetverikov et al. \(2018\)](#), [Cho and White \(2018\)](#), [Christensen and Connault \(2019\)](#), [Fang and Santos \(2019\)](#), [Cattaneo et al. \(2020\)](#) and [Masten and Poirier \(2020\)](#).

When distributions are only partially identified by bounds, the situation is more challenging. The current state of the literature on Makarov bounds focuses on pointwise inference for bound functions (see, e.g., [Fan and Park, 2010, 2012](#); [Fan et al., 2017](#), and [Firpo and Ridder, 2008, 2019](#)). However, the LASD criterion requires a uniform comparison of bound functions, and the map from distribution functions to Makarov bound functions is also not smooth. Fortunately the problem has a similar solution to the point-identified LASD test. The resulting L_2 - and supremum-norm statistics allow us to conduct inference for LASD in the partially identified case using functions that bound the relevant CDFs. The details are included in the online Supplemental Appendix.

We illustrate the practical use of our proposed criteria and tests with a very simple empirical application using data from [Bitler et al. \(2006\)](#). This aims at exemplifying the use of our approach, rather than developing a fully fledged empirical investigation. We show that, in the case of a policy with gainers and losers, the use of our loss aversion-sensitive evaluation criteria may lead to a ranking of policy interventions that differs from that obtained when their outcomes are compared using stochastic dominance.

The rest of the paper is organized as follows. Section 2 presents the basic definitions and notation and defines loss aversion-sensitive dominance. Section 3 develops testable criteria for loss aversion-sensitive dominance. Section 4 proposes statistical inference methods for LASD using sample observations. Section 5 illustrates our methodology using a very simple empirical application that uses data from the experimental evaluation of a well-known welfare policy reform in the US. Section 6 concludes. Our first appendix includes auxiliary results and definitions; our second one collects proof of the results in the paper.

2. Loss aversion-sensitive dominance

In this section, we propose a novel dominance relation for ordering policies under the assumption that social decision makers consider the distribution of individual gains and losses under different policy scenarios. We call this criterion Loss Aversion-Sensitive Dominance (LASD).

Consider a random variable X with cumulative distribution function F . Let $\mathcal{F}$ be the set of cumulative distribution functions with bounded support $\mathcal{X}$. We maintain the assumption throughout that $F \in \mathcal{F}$. The bounded support assumption is made to avoid technical conditions on tails of distribution functions. The aim of this paper is to provide theoretical criteria and econometric methods to rank policy interventions under LASD. The decision maker's goal is to compare policies A and B using the distribution functions of X_A and X_B , labeled F_A and F_B . Because the random variables X_A and X_B represent gains and losses due to the enactment of policies A and B , they represent a change between agents' pre-treatment and post-treatment outcomes. To this end, let Z_0 , Z_A and Z_B represent an agent's potential outcome under the *status quo*, treatment A or treatment B , so that we may write $X_A = Z_A - Z_0$ and $X_B = Z_B - Z_0$. We assume that the variables (Z_0, Z_A, Z_B) have marginal distribution functions (G_0, G_A, G_B) .

It may be assumed that (F_A, F_B) are identified, or that only (G_0, G_A, G_B) are identified. The former case is related to a point-identified model, and the latter to a partially-identified one. Theoretical results for both situations are shown here. Inference for the first situation is considered in this paper, while partially-identified inference results under the second situation are considered in the online Supplemental Appendix.

For example, in the empirical illustration considered in Section 5, we observe quarterly household income before the enactment of a new welfare program, which we assume to represent realizations of Z_0 . Next, we observe income for some households under a continuation of the old program, identifying those as realizations of Z_A , and income for other households under a new welfare program, labeling those observations as Z_B realizations. If policymakers are interested in how one household's earnings evolve over time either by staying with the old program or switching to the new program, then the levels of Z_A and Z_B are not of primary interest, rather the changes represented by X_A and X_B are (and given the longitudinal nature of our data, it is natural to assume that X_A and X_B are identified). Our tests, detailed in Section 4, use the null hypothesis that a household exhibiting loss-aversion would prefer to switch to the new welfare program, and search for evidence to the contrary.

The decision maker has preferences over X (not Z) that are represented via a continuous function.

Definition 2.1 (*Social Value Function (SVF)*). Suppose random variable X has CDF $F \in \mathcal{F}$ and let $W : \mathcal{F} \rightarrow \mathbb{R}$ denote the following social value function

$$W(F) = \int_{\mathcal{X}} v(x) dF(x), \quad (1)$$

where $v : \mathcal{X} \rightarrow \mathbb{R}$ is called a *value function*.³

The social value function defined above is the value assigned to the distribution of X by a social planner that uses the value function v to convert gains and losses into a measure of well-being (Gajdos and Weymark, 2012). This value function v does not have to coincide with any individual's v in the population: as mentioned in the Introduction, the social planner is averse to individual losses either because individuals are loss-averse themselves (political economy motivation), or because she holds normative views that imply her loss aversion towards X . In either case, v will exhibit loss-aversion, i.e. there is asymmetry in the valuation of gains and losses, where losses are weighed more heavily than gains of equal magnitude. Furthermore, v assigns negative value to losses and positive value to gains and is non-decreasing. These properties are formally listed in the next definition.^{4,5}

Definition 2.2 (*Properties of the value function*). The value function $v : \mathcal{X} \rightarrow \mathbb{R}$ is differentiable and satisfies:

1. Disutility of losses and utility of gains: $v(x) \leq 0$ for all $x < 0$, $v(0) = 0$ and $v(x) \geq 0$ for all $x > 0$.
2. Non-decreasing: $v'(x) \geq 0$ for all x .
3. Loss-averse: $v'(-x) \geq v'(x)$ for all $x > 0$.

The properties in Definition 2.2 are typically assumed in Prospect Theory together with the additional requirement of S-shapedness of value function, which we do not consider (see, e.g., p. 279 of Kahneman and Tversky, 1979). Assumptions 1 and 2 are standard monotone increasing conditions. Assumption 3 expresses the idea that “losses loom larger than corresponding gains” and is a widely accepted definition of loss aversion (Tversky and Kahneman, 1992, p.303). It is a stronger condition than the one considered by Kahneman and Tversky (1979).

The following form of $W(F)$ will be useful in subsequent definitions and results.

Proposition 2.3. Suppose that $F \in \mathcal{F}$ and v is differentiable. Then

$$W(F) = - \int_{x \in \mathcal{X} : x \leq 0} v'(x) F(x) dx + \int_{x \in \mathcal{X} : x > 0} v'(x) (1 - F(x)) dx. \quad (2)$$

Assume that the decision maker's social value function W depends on v which satisfies Definition 2.2, and she wishes to compare random variables X_A and X_B which represent gains and losses under two policies labeled A and B . The decision maker prefers X_A over X_B if she evaluates F_A as better than F_B using her SVF — specifically, X_A is preferred to X_B if and only if $W(F_A) \geq W(F_B)$, where W is defined in Definition 2.1. Please note that X_A is preferred to X_B for every v that is described by Definition 2.2. This is what makes dominance conditions robust criteria for comparing distributions. This idea is formalized below.

Definition 2.4 (*Loss Aversion-Sensitive Dominance*). Let X_A and X_B have distribution functions respectively labeled $F_A, F_B \in \mathcal{F}$. If $W(F_A) \geq W(F_B)$ for all value functions v that satisfy Definition 2.2, we say that F_A dominates F_B in terms of *Loss Aversion-Sensitive Dominance*, or LASD for short, and we write $F_A \succeq_{LASD} F_B$.

In the next section we relate this theoretical definition to a more concrete condition that depends on the cumulative distribution functions of the outcome distributions, F_A and F_B .

³ Formally speaking we have $W_v(F)$ but we suppress the subscript v for expositional brevity.

⁴ This standard interpretation of the social welfare function can be further extended. For example, individuals may be uncertain about their counterfactual outcome and form an expectation of $v(\cdot)$ given z_0 . Then we write $W(F) = \iint v(x, z_0) dF_{X|Z_0}(x|z_0) dF_{Z_0}(z_0)$. Denoting a new value function $v^*(z_0) = \int v(x, z_0) dF_{X|Z_0}(x|z_0)$, i.e. an expected value for a given z_0 , $W(F)$ is as in Definition 2.1.

⁵ An interesting direction for future research is to axiomatize the class of social value functions and possibly develop measures of loss aversion based on this class. Some inspiration for axiomatization may come from the inequality and poverty measurement literature. For example, ratio scale invariance (i.e. proportional changes to the units in which gains and losses are measured do not matter), may be a powerful axiom in obtaining a specific functional form.

3. Testable criteria for loss aversion-sensitive dominance

In this section we formulate testable conditions for evaluating distributions of gains and losses in practice. We propose criteria that indicate whether one distribution of gains and losses dominates another in the sense described in [Definition 2.4](#).

Recall that Z_0 , Z_A and Z_B represent an outcome before or after a policy takes effect, while X_A and X_B represent a change from a pre-policy state to an outcome under a policy. The challenge of comparing variables X_A and X_B is well known in the treatment effects literature: because X_A and X_B are defined by differences between the Z_k , F_A and F_B depend on the joint distribution of (Z_0, Z_A, Z_B) , which may not be observable without restrictions imposed by an economic model. In [Section 3.1](#) we abstract from specific identification conditions and discuss LASD under the assumption that F_A and F_B are identified. In [Section 3.2](#) we work with a partially identified case where only the marginal distribution functions G_0 , G_A and G_B are identified and no restrictions are made to identify F_A and F_B .

3.1. The case of point-identified distributions

The LASD concept in [Definition 2.4](#) requires that one distribution is preferred to another over a class of social value functions and is difficult to test directly. The following result relates the LASD concept to a criterion which depends only on marginal distribution functions and orders F_A and F_B according to the class of SVFs allowed in [Definition 2.2](#). In this section we assume that $F_A, F_B \in \mathcal{F}$ are point identified. This may result from a variety of econometric restrictions that deliver identification and are the subject of a large literature.

Theorem 3.1. *Suppose that $F_A, F_B \in \mathcal{F}$. The following are equivalent:*

1. $F_A \succeq_{LASD} F_B$.
2. For all $x \geq 0$, F_A and F_B satisfy

$$F_B(-x) - F_A(-x) \geq \max\{0, F_A(x) - F_B(x)\}. \quad (3)$$

3. For all $x \geq 0$, F_A and F_B simultaneously satisfy

$$F_A(-x) - F_B(-x) \leq 0 \quad (4)$$

and

$$(1 - F_A(x)) - F_A(-x) \geq (1 - F_B(x)) - F_B(-x). \quad (5)$$

[Theorem 3.1](#) provides two different conditions that can be used to verify whether one distribution of gains and losses dominates the other in the LASD sense.⁶ These criteria compare the outcome distributions by examining how the distribution functions (F_A, F_B) assign probabilities to gains and losses of all possible magnitudes. The particular way that they make a comparison is related to the relative importance of gains and losses. Consider condition (3). For the distribution of X_B to be dominated, its distribution function must lie above the distribution of X_A for losses. X_B can be dominated by X_A in the LASD sense even when gains under X_A do not dominate X_B for gains — that is, when $F_A(x) - F_B(x) \geq 0$ for some $x \geq 0$ — as long as this lack of dominance in gains is compensated by sufficient dominance of X_A over X_B in the losses region. This is a consequence of the asymmetric treatment of gains and losses. Conditions (4) and (5) jointly express the same idea, but they help to understand how gains and losses are treated asymmetrically in condition (3). In the losses region, condition (4) is a standard FOSD condition. This is a consequence of loss aversion; note that in the extreme case where only losses matter, we would have (4). In the gains region, dominance has to be sufficiently large so that under X_A , the probability of gains minus the probability of losses (of magnitude x or larger) is no smaller than the corresponding difference for X_B .⁷ Inequality (3) combines the two inequalities represented by (4) and (5) into a single equation.

It is interesting to note that LASD has one property in common with FOSD, namely, a higher mean is a necessary condition for both types of dominance. This follows directly from [Definitions 2.1](#) and [2.2](#) by using $v(x) = x$.

Lemma 3.2. *If $F_A \succeq_{LASD} F_B$ then $E[X_A] \geq E[X_B]$.*

Note that FOSD cannot rank two distributions that have the same mean — that is, if $F_A \succeq_{FOSD} F_B$ and $E[X_A] = E[X_B]$, then $F_A = F_B$. This is not the case for LASD, as the next example demonstrates. Therefore, for example, [Eq. \(3\)](#) may still be used to differentiate between two distributions with the same average effect.

⁶ LASD is a partial order. Over losses, (4) is a partial order because FOSD is a partial order. For the tail condition (5) checking transitivity we have $(1 - F_A(x)) - F_A(-x) \geq (1 - F_B(x)) - F_B(-x)$, $(1 - F_B(x)) - F_B(-x) \geq (1 - F_C(x)) - F_C(-x)$, and $(1 - F_A(x)) - F_A(-x) \geq (1 - F_C(x)) - F_C(-x)$. If $F_A(-x) - F_B(-x) = 0$ then $F_A(-x) = F_B(-x)$ and using it in (5) gives anti-symmetry.

⁷ We leave 1s on both sides of inequality (5) for this interpretation to be more evident.

Example 3.3. Consider the family of uniform distributions on $[-1-y, -y] \cup [y, y+1]$ indexed by $y > 0$ and denote the corresponding member distribution functions F_y . The family of such distributions have mean zero and $F_y \succeq_{LASD} F_{y'}$ whenever $y < y'$. Indeed, note that

$$W(F_y) = \frac{1}{2} \left(\int_{-1-y}^{-y} v(z) dz + \int_y^{1+y} v(z) dz \right)$$

and thus for any v which is loss-averse (see Definition 2.2) we have

$$\begin{aligned} \frac{d}{dy} W(F_y) &= \frac{1}{2} (v(-1-y) - v(-y) + v(1+y) - v(y)) \\ &= - \int_{-1-y}^{-y} v'(z) dz + \int_y^{1+y} v'(z) dz \\ &= \int_y^{1+y} (v'(z) - v'(-z)) dz \leq 0. \end{aligned}$$

It is important to note that LASD is a concept that is specialized to the comparison of distributions that represent gains and losses. Standard FOSD is typically applied to the distribution of outcomes in levels without regard to whether the outcomes resulted from gains or losses of agents relative to a pre-policy state — in our notation, G_A and G_B are typically compared with FOSD, instead of F_A and F_B . FOSD applied to post-policy levels may or may not coincide with LASD applied to changes. This means that even when a strong condition such as FOSD holds for final outcomes, if one took into account how agents value gains and losses it may turn out that the dominant distribution is no longer a preferred outcome. One could apply the FOSD rule to compare distributions of income changes, which implies LASD applied to changes, because FOSD applies to a broader class of value functions. However, this type of comparison would ignore agents' loss aversion, the important qualitative feature that LASD accounts for. The following example shows that the analysis of outcomes in levels using FOSD need not correspond to any LASD ordering of outcomes in changes.

Example 3.4. Let Z_0 represent outcomes before policies A or B . Suppose Z_0 is distributed uniformly over $\{0, 1, 2, 3\}$. Policy A assigns post-policy outcomes depending on the realized Z_0 according to the schedule

$$Z_A = \begin{cases} 3 & \text{if } \{Z_0 = 0\} \\ 2 & \text{if } \{Z_0 = 1\} \\ 0 & \text{if } \{Z_0 = 2\} \\ 1 & \text{if } \{Z_0 = 3\}. \end{cases}$$

Therefore the distribution of $X_A = Z_A - Z_0$ is $P\{X_A = -2\} = 1/2$, $P\{X_A = 1\} = P\{X_A = 3\} = 1/4$. Meanwhile, policy B maintains the status quo: $X_B = Z_B - Z_0 = 0$ with probability 1.

Distributions G_A and G_B are the same, thus FOSD holds between Z_A and Z_B . However, there is no loss aversion-sensitive dominance between X_A and X_B . Indeed, we can find two value functions that fulfill the conditions of Definition 2.4 but order X_A and X_B differently. For example, take $v_1(x) = x^3$. Then $3 = \int v_1(x) dF_A(x) > \int v_1(x) dF_B(x) = 0$. Next let $v_2(x) = \text{sgn}(x)|x|^{1/3}$. Then $-0.02 \approx \int v_2(x) dF_A(x) < \int v_2(x) dF_B(x) \approx 0$.

In the previous example, policy B left pre-treatment outcomes unchanged, or in other words, maintained a *status quo* condition — we had $X_B = Z_B - Z_0 \equiv 0$. Suppose generally that X_B has a distribution that is degenerate at 0. Then $F_B(x) = 0$ for all $x < 0$ and $F_B(x) = 1$ for all $x \geq 0$. We define this as a *status quo* policy distribution, labeled F_{SQ} . When comparison is between a distribution F_A and F_{SQ} , LASD and standard FOSD are equivalent. The distribution that dominates F_{SQ} is necessarily only gains.

Corollary 3.5. Suppose that $F_A \in \mathcal{F}$ and $F_B = F_{SQ}$. Then $F_A \succeq_{LASD} F_{SQ} \iff F_A \succeq_{FOSD} F_{SQ}$.

3.2. The case of partially-identified distributions

In many situations of interest the cumulative distribution functions of gains and losses, F_A and F_B , are not point identified without a model of the relationship between X_A and X_B . However, the marginal distributions of outcomes in levels under different policies, represented by the variables Z_0 , Z_A and Z_B , may be identified. Without information on the dependence between potential outcomes, we can still make some more circumscribed statements with regard to dominance based on bounds for the distribution functions. This section studies the LASD dominance criterion to the case that distribution functions F_A and F_B are only partially identified.

A number of authors have considered functions that bound the distribution functions F_A and F_B . Taking X_A as an example, *Makarov bounds* (Makarov, 1982; Rüschendorf, 1982; Frank et al., 1987) are two functions L and U that satisfy $L(x) \leq F_A(x) \leq U(x)$ for all $x \in \mathbb{R}$, depend only on the marginal distribution functions G_0 and G_A and are pointwise sharp — for any fixed x there exist some Z_0^* and Z_A^* such that the resulting $X_A^* = Z_A^* - Z_0^*$ has a distribution function at x that is equal one of $L(x)$ or $U(x)$. Williamson and Downs (1990) provide convenient definitions for these bound functions. For any two distribution functions G_1, G_2 , define

$$L(x, G_1, G_2) = \sup_{u \in \mathbb{R}} (G_2(u) - G_1(u - x))$$

$$U(x, G_1, G_2) = \inf_{u \in \mathbb{R}} (1 + G_2(u) - G_1(u - x)).$$

For convenience define the policy-specific bound functions for F_k , $k \in \{A, B\}$ and all $x \in \mathbb{R}$, which depend on the marginal CDFs G_0 and G_k , by

$$\begin{aligned} L_k(x) &= L(x, G_0, G_k) \\ &= \sup_{u \in \mathcal{X}} (G_k(u) - G_0(u - x)) \end{aligned} \quad (6)$$

$$\begin{aligned} U_k(x) &= U(x, G_0, G_k). \\ &= 1 + \inf_{u \in \mathcal{X}} (G_k(u) - G_0(u - x)) \end{aligned} \quad (7)$$

Using these definitions we obtain a sufficient and a necessary condition for LASD when only bound functions of the treatment effects distribution functions are observable. The next theorem formalizes the result.

Theorem 3.6. Suppose that $G_0, G_A, G_B \in \mathcal{F}$ and define the bound functions using formulas (6) and (7) for $k \in \{A, B\}$.

1. If for all $x \geq 0$,

$$L_B(-x) - U_A(-x) \geq \max\{0, U_A(x) - L_B(x)\}, \quad (8)$$

then $F_A \succeq_{LASD} F_B$ holds.

2. If $F_A \succeq_{LASD} F_B$ holds, then for all $x \geq 0$

$$U_B(-x) - L_A(-x) \geq \max\{0, L_A(x) - U_B(x)\}. \quad (9)$$

Theorem 3.6 is an analog of **Theorem 3.1** and shows what effect the loss of point identification has on the relationship between dominance and conditions on the CDFs. In particular, one loses a simple “if and only if” characterization that depends on CDFs. Instead, LASD implies a necessary condition using some bound functions, while a different sufficient condition using other bound functions implies LASD. Inference using the necessary condition shown in **Theorem 3.6** is discussed in the online Supplemental Appendix. We remark that there may exist other features of the joint data distribution that do not depend only on pointwise features of the CDFs of changes and would result in a necessary and sufficient condition for LASD under partial identification. That is an interesting open question but is beyond the scope of this paper.

When the comparison is with the *status quo* distribution, the partially identified conditions simplify. **Corollary 3.7** below shows what can be learned about LASD from bound functions in the partially identified case.

Corollary 3.7. Suppose that $F_B = F_{SQ}$ and that $G_0, G_A \in \mathcal{F}$. Define the bound functions U_A and L_A using formulas (6) and (7). Then $U_A(-x) = 0$ for all $x \geq 0 \Rightarrow F_A \succeq_{LASD} F_{SQ}$ and $F_A \succeq_{LASD} F_{SQ} \Rightarrow L_A(-x) = 0$ for all $x \geq 0$.

4. Inferring loss aversion-sensitive dominance

In this section we propose statistical inference methods for the loss aversion-sensitive dominance (LASD) criterion discussed in previous sections. We consider the null and alternative hypotheses

$$\begin{aligned} H_0 &: F_A \succeq_{LASD} F_B \\ H_1 &: F_A \not\succeq_{LASD} F_B. \end{aligned} \quad (10)$$

Under the null hypothesis (10) policy A dominates B in the LASD sense, similar to much of the literature on stochastic dominance. It is a simplification of the hypotheses considered for several potential policies discussed in [Linton et al. \(2005\)](#), who test whether one policy is maximal, and the techniques developed below could be extended to compare several policies in the same way in a straightforward manner.⁸ The null hypothesis above represents the assumption that policy A is preferred by agents in the LASD sense. Rejection of the null implies that there is significant evidence for ambiguity in the ordering of the policies by LASD. Unfortunately, a drawback of the proposed procedure is that rejection of the null does not inform one about which sort of value function v results in a rejection. Strong orderings of policies can result in more information, although they constrain v by construction, and such exploration is left for future research.⁹

We consider tests for this null hypothesis given sample data observed under two different identification assumptions. We start with the case where one can directly observe samples $\{X_{Ai}\}_{i=1}^{n_A}$ and $\{X_{Bi}\}_{i=1}^{n_B}$ which represent agents’ gains and losses, or in other

⁸ [Linton et al. \(2005\)](#) consider a test for Prospect Theory by testing whether the integral of one CDF dominates the other. This paper considers a different approach in which we impose loss aversion on the value function, and then derive testable conditions on the CDFs.

⁹ There is also another strand of literature that develops methods to estimate the optimal treatment assignment policy that maximizes a social welfare function. Recent developments can be found in [Manski \(2004\)](#), [Dehejia \(2005\)](#), [Hirano and Porter \(2009\)](#), [Stoye \(2009\)](#), [Bhattacharya and Dupas \(2012\)](#), [Tetenov \(2012\)](#), [Kitagawa and Tetenov \(2018, 2021\)](#), among others. These papers focus on the decision-theoretic properties and procedures that map empirical data into treatment choices. In this literature, our paper is most closely related to [Kasy \(2016\)](#), which focuses on welfare rankings of policies rather than optimal policy choice.

words, we simply assume that the distribution functions of X_A and X_B are point-identified and their distribution functions can be estimated using the empirical distribution functions from two samples. Next we extend these results to the partially-identified case where no assumption about the joint distribution of potential outcomes under either treatment is made. In this case, we assume that three samples are observable, $\{Z_{0i}\}_{i=1}^{n_0}$, $\{Z_{Ai}\}_{i=1}^{n_A}$ and $\{Z_{Bi}\}_{i=1}^{n_B}$, representing outcomes under a control or pre-policy state and outcomes under policies A and B . Then tests are based on plug-in estimates for bounds for $X_A = Z_A - Z_0$ and $X_B = Z_B - Z_0$.

We consider distribution functions as members of the space of bounded functions on the support $\mathcal{X} \subseteq \mathbb{R}$, denoted $\ell^\infty(\mathcal{X})$, equipped with the supremum norm, defined for $g : \mathbb{R}^k \rightarrow \mathbb{R}^\ell$ by $\|g\|_\infty = \max_j \{\sup_{x \in \mathbb{R}^k} |g_j(x)|\}$. For real numbers x let $(x)^+ = \max\{0, x\}$. Given a sequence of bounded functions $\{g_n\}_n$ and limiting random element g we write $g_n \rightsquigarrow g$ to denote weak convergence in $(\ell^\infty, \|\cdot\|_\infty)$ in the sense of Hoffman-Jørgensen (van der Vaart and Wellner, 1996).

4.1. Inferring dominance from point identified treatment distributions

In this subsection we suppose that the pair of marginal distribution functions $F = (F_A, F_B)$ is identified. In the Online Supplemental Appendix C, we provide results extending the dominance tests to the case that distribution functions F_A and F_B are only partially identified.

4.1.1. Test statistics

To implement a test of the hypotheses (10) we employ the results of Theorem 3.1 to construct maps of F into criterion functions that are used to detect deviations from the hypothesis H_0 . Specifically, recalling that $(x)^+ = \max\{0, x\}$, for the point-identified case we examine maps $T_1 : (\ell^\infty(\mathbb{R}))^2 \rightarrow \ell^\infty(\mathbb{R}_+)$ and $T_2 : (\ell^\infty(\mathbb{R}))^2 \rightarrow (\ell^\infty(\mathbb{R}_+))^2$, defined for each $x \geq 0$ by

$$T_1(F)(x) = (F_A(x) - F_B(x))^+ + F_A(-x) - F_B(-x) \quad (11)$$

and

$$T_2(F)(x) = \begin{bmatrix} F_A(-x) - F_B(-x) \\ F_A(x) - F_B(x) + F_A(-x) - F_B(-x) \end{bmatrix}. \quad (12)$$

Functions $T_1(F)$ and $T_2(F)$ are designed so that large positive values will indicate a violation of the null. Taking T_1 as an example, Theorem 3.1 states that $W(F_A) \geq W(F_B)$ if and only if $F_B(-x) - F_A(-x) \geq (F_A(x) - F_B(x))^+$ for all $x \geq 0$, so tests can be constructed by looking for x where $T_1(F)(x)$ becomes significantly positive. We will refer to T_j as maps from pairs of distribution functions to another function space, and also refer to them as functions.

The hypotheses (10) can be rewritten in two equivalent forms, depending on whether one uses T_1 or T_2 to transform distribution functions: letting $\mathcal{X} \subseteq \mathbb{R}_+$ be an evaluation set, we have

$$\begin{aligned} H_0^{(1)} : T_1(F)(x) &\leq 0, \quad \text{for all } x \in \mathcal{X}, \\ H_1^{(1)} : T_1(F)(x) &> 0, \quad \text{for some } x \in \mathcal{X} \end{aligned} \quad (13)$$

and

$$\begin{aligned} H_0^{(2)} : T_2(F)(x) &\leq 0_2, \quad \text{for all } x \in \mathcal{X}, \\ H_1^{(2)} : T_2(F)(x) &\not\leq 0_2, \quad \text{for some } x \in \mathcal{X}. \end{aligned} \quad (14)$$

In the second set of hypotheses 0_2 is a two-dimensional vector of zeros and inequalities are taken coordinate-wise.

The next step in testing the hypotheses (13) and (14) is to estimate $T_1(F)$ and $T_2(F)$. Let $\mathbb{F}_n = (\mathbb{F}_{An}, \mathbb{F}_{Bn})$ denote the pair of marginal empirical distribution functions, that is, $\mathbb{F}_{kn}(x) = \frac{1}{n_k} \sum_{i=1}^{n_k} \mathbf{1}\{X_{ki} \leq x\}$ for $k \in \{A, B\}$. These are well-behaved estimators of the components of F . Letting $n = n_A + n_B$, standard empirical process theory shows that $\sqrt{n}(\mathbb{F}_n - F)$ converges weakly to a Gaussian process under weak assumptions (van der Vaart, 1998, Example 19.6). In order to conduct inference for loss aversion-sensitive dominance, we use plug-in estimators $T_j(\mathbb{F}_n)$ for $j \in \{1, 2\}$. See Remark B.1 in Appendix D for details on the computation of these functions.

In order to detect when $T_j(\mathbb{F}_n)$ is significantly positive, we consider statistics based on a one-sided supremum norm or a one-sided L_2 norm over $\mathcal{X}$. Kolmogorov–Smirnov (i.e., supremum norm) type statistics are

$$V_{1n} = \sqrt{n} \sup_{x \in \mathcal{X}} (T_1(\mathbb{F}_n)(x))^+ \quad (15)$$

$$V_{2n} = \sqrt{n} \max \left\{ \sup_{x \in \mathcal{X}} (T_{21}(\mathbb{F}_n)(x))^+, \sup_{x \in \mathcal{X}} (T_{22}(\mathbb{F}_n)(x))^+ \right\}. \quad (16)$$

Meanwhile Cramér-von Mises (or L_2 norm) test statistics are defined by

$$W_{1n} = \sqrt{n} \left(\int_{\mathcal{X}} ((T_1(\mathbb{F}_n)(x))^+)^2 dx \right)^{1/2}, \quad (17)$$

$$W_{2n} = \sqrt{n} \left(\int_{\mathcal{X}} ((T_{21}(\mathbb{F}_n)(x))^+)^2 + ((T_{22}(\mathbb{F}_n)(x))^+)^2 dx \right)^{1/2}. \quad (18)$$

Because all the CDFs used in these statistics belong to $\mathcal{F}$, distributions with bounded support, the integrands in the L_2 statistics are square-integrable.

4.1.2. Limiting distributions

We wish to establish the limiting distributions of V_{jn} and W_{jn} , for $j \in \{1, 2\}$, under the null hypothesis $H_0 : F_A \succeq_{LASD} F_B$. Two challenges arise when considering these test statistics. First, the form of the null hypothesis as a functional inequality to be tested uniformly over $\mathcal{X}$ is a source of irregularity. Let the joint probability distribution of (X_A, X_B) be denoted by P . Because the null hypothesis, $F_A \succeq_{LASD} F_B$, is a functional weak inequality the asymptotic distributions of the test statistics V_j and W_j may depend on features of P . This is referred to as *non-uniformity in P* in Linton et al. (2010), Andrews and Shi (2013), and requires attention when resampling.

Second, due to the pointwise maximum function in its definition, T_1 is too irregular as a map from the data to the space of bounded functions to establish a limiting distribution for the empirical process $\sqrt{n}(T_1(\mathbb{F}_n) - T_1(F))$ using conventional statistical techniques. In contrast, T_2 is a linear map of F , which implies that $\sqrt{n}(T_2(\mathbb{F}_n) - T_2(F))$ has a well-behaved limiting distribution in $(\ell^\infty(\mathbb{R}_+))^2$.¹⁰

Despite the above challenges, we show that V_{jn} and W_{jn} (for $j \in \{1, 2\}$) have well-behaved asymptotic distributions, and furthermore, that the limiting random variables satisfy $V_1 \sim V_2$ and $W_1 \sim W_2$. This is an important result because it is the foundation for applying bootstrap techniques for inference. Before stating the formal assumptions and asymptotic properties of the tests, we discuss the two difficulties mentioned above in more detail.

The limiting distributions of V_{jn} and W_{jn} statistics depend on features of P . Let $\mathcal{P}_0$ be the set of distributions P such that $F_A \succeq_{LASD} F_B$. These are distributions with marginal distribution functions F such that $T_j(F)(x) \leq 0$ for all $x \geq 0$. To discuss the relationship between these sets of distributions and test statistics, we relabel the two coordinates of the T_2 function as

$$m_1(x) = F_A(-x) - F_B(-x) \quad (19)$$

and

$$m_2(x) = F_A(-x) - F_B(-x) + F_A(x) - F_B(x). \quad (20)$$

When $P \in \mathcal{P}_0$, both $m_1(x) \leq 0$ and $m_2(x) \leq 0$ for all $x \geq 0$.

More detail is required about the behavior of the two coordinate functions to determine the limiting distributions of V_{jn} and W_{jn} statistics. For L_2 -norm statistics W_{1n} and W_{2n} , we define the following relevant subdomains of $\mathcal{X}$, which collect the arguments in the interior of $\mathcal{X}$ where m_1 or m_2 are equal to zero:

$$\mathcal{X}_0^1(P) = \{x \in \text{int}\mathcal{X} : m_1(x) = 0\} \quad (21)$$

$$\mathcal{X}_0^2(P) = \{x \in \text{int}\mathcal{X} : m_2(x) = 0\}. \quad (22)$$

Denote $\mathcal{X}_0(P) \subseteq \mathcal{X}$ as the set of x where $T_1(F)(x) = 0$ or at least one coordinate of $T_2(F)$ equals 0 for probability distribution P . As will be seen below, $\mathcal{X}_0(P)$ is the same for both the T_1 and T_2 functions, and when it is non-empty, test statistics have a nondegenerate distribution. Following Linton et al. (2010), we call $\mathcal{X}_0(P)$ the contact set for the distribution P . Given the above definitions, under the null hypothesis we can write

$$\mathcal{X}_0(P) = \mathcal{X}_0^1(P) \cup \mathcal{X}_0^2(P).$$

On the other hand, the supremum-norm statistics V_{1n} and V_{2n} need a different family of sets, namely the sets of ϵ -maximizers of m_1 and m_2 . For any $\epsilon \geq 0$ and $k \in \{1, 2\}$, let

$$\mathcal{M}^k(\epsilon) = \left\{ x \in \mathcal{X} : m_k(x) \geq \sup_{x \in \mathcal{X}} m_k(x) - \epsilon \right\}. \quad (23)$$

An important subset of $\mathcal{P}_0$ are those P for which test statistics have nontrivial limiting distributions under the null hypothesis — that is, not degenerate at 0, which occurs when there is some x such that $T_j(F)(x) = 0$ (note that there are no x such that $T_j(F)(x) > 0$ when $P \in \mathcal{P}_0$). Define $\mathcal{P}_{00} \subset \mathcal{P}_0$ to be the set of all P such that $\mathcal{X}_0(P) \neq \emptyset$. If $P \in \mathcal{P}_0 \setminus \mathcal{P}_{00}$ then $\mathcal{X}_0(P) = \emptyset$ and because the distribution satisfies the null hypothesis, F_A strictly dominates F_B everywhere and the criterion functions T_j are strictly negative over $\mathcal{X}$. When $P \in \mathcal{P}_0 \setminus \mathcal{P}_{00}$, test statistics have asymptotic distributions that are degenerate at zero because test statistics will detect that policy A is strictly better than B over all of $\mathcal{X}$. When $P \in \mathcal{P}_{00}$, $T_j(F)$ is zero over $\mathcal{X}_0(P)$ and test statistics have a nontrivial asymptotic distribution over $\mathcal{X}_0(P)$. Thus, when $F_A \succeq_{LASD} F_B$, the asymptotic behavior of test statistics depends on whether $P \in \mathcal{P}_{00}$ or $P \in \mathcal{P}_0 \setminus \mathcal{P}_{00}$. Note that when $P \in \mathcal{P}_{00}$, we have $\lim_{\epsilon \searrow 0} \mathcal{M}^k(\epsilon) = \mathcal{X}_0^k(P)$ (that is, nonstochastic convergence in the sense of Painlevé-Kuratowski, see, e.g., Rockafellar and Wets (1998, p. 111)) for whichever coordinate function actually achieves the maximal value zero.

The second challenge for testing is related to the scaled difference $\sqrt{n}(T_1(\mathbb{F}_n) - T_1(F))$ as n grows large. Hadamard differentiability is an analytic tool used to establish the asymptotic distribution of nonlinear maps of the empirical process. Definition D.1 in Appendix D provides a precise statement of the concept. When a map is Hadamard differentiable — for example T_2 , which is linear as a map from $(\ell^\infty(\mathbb{R}))^2$ to $(\ell^\infty(\mathbb{R}_+))^2$ and is thus trivially differentiable — the functional delta method can be applied to describe its asymptotic behavior as a transformed empirical process, and a chain rule makes the analysis of compositions of several Hadamard-differentiable maps tractable. Also, the Hadamard differentiability of a map implies resampling is consistent when this map is applied

¹⁰ The issues of a general lack of differentiability of functions arrived at by marginal optimization and a solution for inference based on directly characterizing the behavior of test statistics applied to such functions are studied in more generality in Firpo et al. (2023). However, we highlight that the tests described here are extensions of the results of that paper and are specifically tailored to this application.

to the resampled empirical process (van der Vaart, 1998, Theorem 23.9) — so, for example, the distribution of resampled criterion processes $\sqrt{n}(T_2(\mathbb{F}_n^*) - T_2(\mathbb{F}_n))$ is a consistent estimate of the asymptotic distribution of $\sqrt{n}(T_2(\mathbb{F}_n) - T_2(F))$ in the space $\ell^\infty(\mathbb{R}_+)$. On the other hand, consider the T_1 map. The pointwise Hadamard directional derivative of $T_1(f)(x)$ at a given $x \geq 0$ in direction $h(x) = (h_A(x), h_B(x))$ is

$$T'_{1f}(h)(x) = \begin{cases} h_A(x) - h_B(x) + h_A(-x) - h_B(-x), & f_A(x) > f_B(x) \\ (h_A(x) - h_B(x))^+ + h_A(-x) - h_B(-x), & f_A(x) = f_B(x) \\ h_A(-x) - h_B(-x), & f_A(x) < f_B(x). \end{cases} \quad (24)$$

This map, thought of as a map between function spaces, $(\ell^\infty(\mathbb{R}))^2$ and $\ell^\infty(\mathbb{R}_+)$, is not differentiable because the scaled differences $(T_1(f)(x) - T_1(f + th_t)(x))/t$ converge to the above derivative at each point x , but may not converge uniformly in $\mathbb{R}_+$. Despite the lack of differentiability of the map $F \mapsto T_1(F)$, we show in Lemma D.3 in Appendix D that the maps $F \mapsto V_1$ and $F \mapsto W_1$ are Hadamard directionally differentiable, which implies these maps are just regular enough that existing statistical methods can be applied to their analysis. Later in this section we apply the resampling technique recently developed in Fang and Santos (2019) along with this directional differentiability to describe hypothesis tests using V_{1n} or W_{1n} .

Having discussed the difficulties in the relationship between distributions and test statistics, we turn to assumptions on the observations. In order to conduct inference using either $T_1(\mathbb{F}_n)$ or $T_2(\mathbb{F}_n)$ we make the following assumptions.

- A1 The observations $\{X_{Ai}\}_{i=1}^{n_A}$ and $\{X_{Bi}\}_{i=1}^{n_B}$ are iid samples and independent of each other and are continuously distributed with marginal distribution functions F_A and F_B respectively.
- A2 Let the sample sizes n_A and n_B increase in such a way that $n_k/(n_A + n_B) \rightarrow \lambda_k$ as $n_A, n_B \rightarrow \infty$, where $0 < \lambda_k < 1$ for $k \in \{A, B\}$. Define $n = n_A + n_B$.

Under these assumptions we establish the asymptotic properties of the test statistics under the null and fixed alternatives. Under the above assumptions, there is a Gaussian process $\mathcal{G}_F$ such that $\sqrt{n}(\mathbb{F}_n - F) \rightsquigarrow \mathcal{G}_F$. We denote each coordinate process $\mathcal{G}_{F_A}$ and $\mathcal{G}_{F_B}$, and for convenience define two transformed processes: for each $x \geq 0$ let

$$\mathcal{G}_1(x) = \mathcal{G}_{F_A}(-x) - \mathcal{G}_{F_B}(-x) \quad (25)$$

$$\mathcal{G}_2(x) = \mathcal{G}_{F_A}(x) - \mathcal{G}_{F_B}(x) - \mathcal{G}_{F_A}(-x) + \mathcal{G}_{F_B}(-x). \quad (26)$$

These will be used in the theorem below.

Theorem 4.1. Make Assumptions A1–A2. Define the limiting Gaussian processes $\mathcal{G}_1$ and $\mathcal{G}_2$ as above. Then:

1. Suppose that $P \in \mathcal{P}_{00}$. As $n \rightarrow \infty$, $V_{1n} \rightsquigarrow V_1$ and $W_{1n} \rightsquigarrow W_1$, where

$$V_1 \sim \max \left\{ 0, \sup_{x \in \mathcal{X}_0^1(P)} \mathcal{G}_1(x) \cdot \mathbf{1} \left\{ \sup_{x \in \mathcal{X}} m_1(x) = 0 \right\}, \sup_{x \in \mathcal{X}_0^2(P)} \mathcal{G}_2(x) \cdot \mathbf{1} \left\{ \sup_{x \in \mathcal{X}} m_2(x) = 0 \right\} \right\}$$

and

$$W_1 \sim \left(\int_{\mathcal{X}_0^1(P)} \left((\mathcal{G}_1(x))^+ \right)^2 dx + \int_{\mathcal{X}_0^2(P)} \left((\mathcal{G}_2(x))^+ \right)^2 dx \right)^{1/2}.$$

2. Suppose that $P \in \mathcal{P}_{00}$. As $n \rightarrow \infty$, $V_{2n} \rightsquigarrow V_2$ and $W_{2n} \rightsquigarrow W_2$, where $V_2 \sim V_1$ and $W_2 \sim W_1$.
3. Suppose that $P \in \mathcal{P}_0 \setminus \mathcal{P}_{00}$ for $j = 1$ or 2 . As $n \rightarrow \infty$, $P\{V_{jn} > \epsilon\} \rightarrow 0$ and $P\{W_{jn} > \epsilon\} \rightarrow 0$ for all $\epsilon > 0$.
4. Suppose that $P \notin \mathcal{P}_0$. As $n \rightarrow \infty$, $P\{V_{jn} > c\} \rightarrow 1$ and $P\{W_{jn} > c\} \rightarrow 1$ for all $c \geq 0$ for $j = 1$ or 2 .

Theorem 4.1 derives the asymptotic properties of the proposed test statistics. Parts 1 and 2 establish the weak limits of V_{jn} and W_{jn} for $j \in \{1, 2\}$ when the null hypothesis is true. Recall that when $P \in \mathcal{P}_{00}$, $\lim_{\epsilon \searrow 0} \mathcal{M}^k(\epsilon) = \mathcal{X}_0^k(P)$, which is why $\mathcal{M}^k(\epsilon)$ terms are absent in the first part of the theorem. Remarkably, the test statistics using T_1 and T_2 criterion processes have the same asymptotic behavior despite the different appearances of the underlying processes and the irregularity of T_1 . Part 3 shows that the statistics are asymptotically degenerate at zero when the contact set is empty, that is, when P lies on the interior of the null region. Part 4 shows that the test statistics diverge when data comes from any distribution that does not satisfy the null hypothesis.

The limiting distributions described in Part 1 of Theorem 4.1 are not standard because the distributions of the test statistics depend on features of P through the $\mathcal{X}_0(P)$ terms in each expression. Therefore, to make practical inference feasible, we suggest the use of resampling techniques below.

4.1.3. Resampling procedures for inference

The proposed test statistics have complex limiting distributions. In this subsection, we present resampling procedures to estimate the limiting distributions of both V_{jn} and W_{jn} for $j \in \{1, 2\}$ under the assumption that $P \in \mathcal{P}_{00}$. Naive use of bootstrap data generating processes in the place of the original empirical process suffers from distortions due to discontinuities in the directional derivatives of the maps that define the distributions of the test statistics. In finite samples the plug-in estimate will not find, for example, the

region where $F_A(x) - F_B(x) = 0$, where the derivatives exhibit discontinuous behavior. Our procedure involves making estimates of the derivatives involved in the limiting distribution and a standard exchangeable bootstrap routine, as proposed in [Fang and Santos \(2019\)](#).¹¹

In order to estimate contact sets, define a sequence of constants $\{a_n\}$ such that $a_n \searrow 0$ and $\sqrt{n}a_n \rightarrow \infty$ and let $\hat{m}_{1n}(x) = \mathbb{F}_{An}(-x) - \mathbb{F}_{Bn}(-x)$ and $\hat{m}_{2n}(x) = \mathbb{F}_{An}(-x) - \mathbb{F}_{Bn}(-x) + \mathbb{F}_{An}(x) - \mathbb{F}_{Bn}(x)$. Then for W_j statistics define estimated contact sets by

$$\hat{\mathcal{X}}_0^1 = \{x \in \text{int}\mathcal{X} : |\hat{m}_{1n}(x)| \leq a_n\} \quad (27)$$

$$\hat{\mathcal{X}}_0^2 = \{x \in \text{int}\mathcal{X} : |\hat{m}_{2n}(x)| \leq a_n\}. \quad (28)$$

When both sets are empty, replace both estimates by $\mathcal{X}$, as suggested in [Linton et al. \(2010\)](#) to ensure nondegenerate bootstrap reference distributions. Meanwhile, for V_j statistics define estimated ϵ -maximizer sets. For a sequence of constants $\{b_n\}$ such that $b_n \searrow 0$ and $\sqrt{n}b_n \rightarrow \infty$, let

$$\hat{\mathcal{M}}^1 = \{x \in \mathcal{X} : \hat{m}_{1n}(x) \geq \max \hat{m}_{1n}(x) - b_n\}, \quad (29)$$

$$\hat{\mathcal{M}}^2 = \{x \in \mathcal{X} : \hat{m}_{2n}(x) \geq \max \hat{m}_{2n}(x) - b_n\}. \quad (30)$$

Although the null hypothesis may imply that the maximum $m_1(x)$ is zero, the above formulas use the maximum of the sample analog without setting its maximum equal to zero, which is important for ensuring non-empty set estimates. Using these estimates, the distributions of V_1 and W_1 can be estimated from sample data (recall that Part 2 of [Theorem 4.1](#) asserts that these are the same distributions as those of V_2 and W_2). We conducted simulation experiments to choose these parameters using a few simulated data-generating processes, which are briefly discussed in the appendix in the context of simulations that suggest that the resulting tests have correct size and good power. Scaling the estimated processes by their pointwise standard deviation functions when estimating contact sets as in [Lee et al. \(2018\)](#) might result in better performance when distribution functions are evaluated near their tails, but we leave that rather complex topic for future research.

Resampling routine to estimate the distributions of V_{jn} and W_{jn} for $j = 1, 2$:

1. If using a Cramér-von Mises statistic, given a sequence of constants $\{a_n\}$, estimate the contact sets $\hat{\mathcal{X}}_0^1$ and $\hat{\mathcal{X}}_0^2$. If using a Kolmogorov-Smirnov statistic, given a sequence of constants $\{b_n\}$, estimate the b_n -maximizer sets of $\hat{m}_{1n}$ and $\hat{m}_{2n}$.

Next repeat the following two steps for $r = 1, \dots, R$:

2. Construct the resampled processes

$$\begin{aligned} F_{r1n}^*(x) &= \sqrt{n} \left(\mathbb{F}_{An}^*(-x) - \mathbb{F}_{Bn}^*(-x) - \mathbb{F}_{An}(-x) + \mathbb{F}_{Bn}(-x) \right) \\ F_{r2n}^*(x) &= \sqrt{n} \left(\mathbb{F}_{An}^*(-x) - \mathbb{F}_{Bn}^*(-x) - \mathbb{F}_{An}(-x) + \mathbb{F}_{Bn}(-x) \right. \\ &\quad \left. + \mathbb{F}_{An}^*(x) - \mathbb{F}_{Bn}^*(x) - \mathbb{F}_{An}(x) + \mathbb{F}_{Bn}(x) \right) \end{aligned}$$

using an exchangeable bootstrap.

3. Calculate the resampled test statistic. Letting $\hat{k} = \arg\max_k \{\sup_{x \geq 0} \hat{m}_{kn}(x)\}$ and $\{c_n\} \searrow 0$ satisfy $\sqrt{n}c_n \rightarrow \infty$, calculate

$$V_{rn}^* = \begin{cases} \left(\max_{x \in \hat{\mathcal{M}}^k} F_{rk n}^*(x) \right)^+ & | \max \hat{m}_{1n} - \max \hat{m}_{2n} | > c_n \\ \max \{0, \max_{x \in \hat{\mathcal{M}}^1} F_{r1n}^*(x), \max_{x \in \hat{\mathcal{M}}^2} F_{r2n}^*(x)\} & | \max \hat{m}_{1n} - \max \hat{m}_{2n} | \leq c_n \end{cases} \quad (31)$$

or

$$W_{rn}^* = \left(\int_{\hat{\mathcal{X}}_0^1} \left((F_{r1n}^*(x))^+ \right)^2 dx + \int_{\hat{\mathcal{X}}_0^2} \left((F_{r2n}^*(x))^+ \right)^2 dx \right)^{1/2}. \quad (32)$$

Finally,

4. Let $\hat{q}_{V^*}(1 - \alpha)$ and $\hat{q}_{W^*}(1 - \alpha)$ be the $(1 - \alpha)$ th sample quantile from the bootstrap distributions of $\{V_{rn}^*\}_{r=1}^R$ or $\{W_{rn}^*\}_{r=1}^R$, respectively, where $\alpha \in (0, 1)$ is the nominal size of the tests. Reject the null hypothesis (13) or (14) if V_{jn} and W_{jn} defined in (15)–(18) are, respectively, larger than $\hat{q}_{V^*}(1 - \alpha)$ or $\hat{q}_{W^*}(1 - \alpha)$.

The formulas in part 3 of the steps above are obtained by inserting estimated contact sets and resampled empirical processes in the place of population-level quantities into the functions shown in part 1 of [Theorem 4.1](#).

The resampled statistics are calculated by imposing the null hypothesis and assuming that the region $\mathcal{X}_0^j(P)$ is the only part of the domain that provides a nondegenerate contribution to the asymptotic distribution of the statistic under the null. The two cases of each part in the maximum arise from trying to impose the null behavior on the resampled supremum norm statistics,

¹¹ Given a set of weights $\{W_i\}_{i=1}^n$ that sum to one and are independent of $\{X_i\}_{i=1}^n$, the exchangeable bootstrap measure is a randomly-weighted measure that puts mass W_i at observed sample point X_i for each i . This encompasses, for example, the standard bootstrap, m -of- n bootstrap and wild bootstrap. See Section 3.6.2 of [van der Vaart and Wellner \(1996\)](#) for more specific details.

even when it appears the null is violated based on the value of the sample statistic. A simple alternative way to conduct inference would be to assume the least-favorable null hypothesis that $F_A \equiv F_B$, and to resample using all of $\mathcal{X}$. However, this may result in tests with lower power (Linton et al., 2010) — power loss arises in situations where $\mathcal{X}_0(P) \subset \mathcal{X}$ (strictly), so that the T_j process is only nondegenerate on a subset, while bootstrapped processes that assume $\mathcal{X}_0(P) = \mathcal{X}$ would look over all of $\mathcal{X}$ and result in a stochastically larger bootstrap distribution than the true distribution.

The next result shows that our tests based on the resampling schemes described above have accurate size under the null hypothesis. In order to metrize weak convergence we use test functions from the set BL_1 , which denotes Lipschitz functions $\mathbb{R} \rightarrow \mathbb{R}$ that have constant 1 and are bounded by 1.

Theorem 4.2. *Make Assumptions A1–A2 and suppose that $P \in \mathcal{P}_{00}$. Let X denote the sample observations. Then for $j = 1, 2$, the bootstrap is consistent:*

$$\sup_{f \in BL_1} \left| \mathbb{E} [f(V_n^*) | X] - \mathbb{E} [f(V_1)] \right| = o_P(1)$$

and

$$\sup_{f \in BL_1} \left| \mathbb{E} [f(W_n^*) | X] - \mathbb{E} [f(W_1)] \right| = o_P(1),$$

where V_1 and W_1 are defined in Theorem 4.1. In particular, when $P \in \mathcal{P}_{00}$ the resampling procedure outlined above results in asymptotically valid inference: for any $P \in \mathcal{P}_{00}$, letting $q_{V_j^*}(1 - \alpha) = \lim_{R \rightarrow \infty} \hat{q}_{V_j^*}(1 - \alpha)$ and $q_{W_j^*}(1 - \alpha) = \lim_{R \rightarrow \infty} \hat{q}_{W_j^*}(1 - \alpha)$,

$$\limsup_{n \rightarrow \infty} P \left\{ V_{jn} > q_{V_j^*}(1 - \alpha) \right\} \leq \alpha$$

and

$$\limsup_{n \rightarrow \infty} P \left\{ W_{jn} > q_{W_j^*}(1 - \alpha) \right\} \leq \alpha,$$

with equality when the distributions of V_j and W_j are strictly increasing at their $(1 - \alpha)$ -th quantiles.

The result in above theorem is stated in terms of the limiting variables V_1 and W_1 and bootstrap analogs. V_1 and W_1 , using the functional delta method, are Hadamard directional derivatives of a chain of maps from the marginal distribution functions F to the real line, and the derivatives are most compactly expressed as the definitions in Theorem 4.1.

The bootstrap variables combine conventional resampling with finite-sample estimates of the maps defined in Part 1 of Theorem 4.1, which is a resampling approach proposed in Fang and Santos (2019). Their result is actually more general — it states that with a more flexible estimator V_n^* , we would obtain bootstrap consistency for P in the null and alternative regions. Because our focus is on testing $F_A \succeq_{LASD} F_B$, however, our resampling scheme, and Theorem 4.2, are done under the imposition of the null hypothesis. The resampling consistency result in Theorem 4.2 implies that our bootstrap tests have asymptotically correct size for all probability distributions in the null region, in the same sense as was stressed in Linton et al. (2010). A formal statement showing size control over all of $\mathcal{P}_0$ is given in Theorem 4.3 in Appendix D. Along with Part 4 of Theorem 4.1, Theorem 4.3 additionally implies that our tests are consistent, that is, that their power to detect violations from the null represented by fixed alternative distributions tends to one. This is because the resampling scheme produces asymptotically bounded critical values, while the test statistics diverge under the alternative.

The behavior of bootstrap tests under the null and alternatives is most easily examined using distributions local to P . We consider sequences of distributions P_n local to the null distribution P such that for a mean-zero, square-integrable function η , P_n have distribution functions F_n (where P has CDF F) that satisfy

$$\lim_{n \rightarrow \infty} \int \left(\sqrt{n} \left(\sqrt{dF_n} - \sqrt{dF} \right) - \frac{1}{2} \eta \sqrt{dF} \right)^2 \rightarrow 0. \quad (33)$$

The behavior of the underlying empirical process under local alternatives satisfies Assumption 5 of Fang and Santos (2019) in a straightforward way (Wellner, 1992, Theorem 1).

Theorem 4.3. *Make Assumptions A1–A2 and suppose that $F_A \succeq_{LASD} F_B$. Suppose that $\mathcal{X}$ is convex. Let $\hat{q}_{V_j^*}(1 - \alpha)$ and $\hat{q}_{W_j^*}(1 - \alpha)$ be the $(1 - \alpha)$ th sample quantile from the bootstrap distributions as described in the routines above, and let $q_{V_j^*}(1 - \alpha) = \lim_{R \rightarrow \infty} \hat{q}_{V_j^*}(1 - \alpha)$ and $q_{W_j^*}(1 - \alpha) = \lim_{R \rightarrow \infty} \hat{q}_{W_j^*}(1 - \alpha)$. Then for $j = 1, 2$,*

1. When $P \in \mathcal{P}_0$ and $\{P_n\}$ satisfy (33) and $T_j(F_n)(x) \leq 0$ for all $x \geq 0$,

$$\limsup_{n \rightarrow \infty} P_n \left\{ V_{jn} > q_{V_j^*}(1 - \alpha) \right\} \leq \alpha$$

and

$$\limsup_{n \rightarrow \infty} P_n \left\{ W_{jn} > q_{W_j^*}(1 - \alpha) \right\} \leq \alpha.$$

2. When $P \in \mathcal{P}_{00}$ and $\{P_n\}$ satisfy (33) and $T_j(F_n)(x) \leq 0$ for all $x \geq 0$, and the distribution of V or W is increasing at its $(1 - \alpha)^{\text{th}}$ quantile,

$$\lim_{n \rightarrow \infty} P_n \left\{ V_{jn} > q_{V^*}(1 - \alpha) \right\} = \alpha$$

and

$$\lim_{n \rightarrow \infty} P_n \left\{ W_{jn} > q_{W^*}(1 - \alpha) \right\} = \alpha.$$

In the Supplemental Appendix we provide Monte Carlo numerical evidence of the finite sample properties of both point- and partially-identified methods. The simulations show that tests have empirical size close to the nominal, and high power against selected alternatives.

5. Empirical illustration

In this section we briefly illustrate the use of our approach using household-level data from a well-known experimental evaluation of alternative welfare programs in the state of Connecticut, documented in [Bitler et al. \(2006\)](#). Aid to Families with Dependent Children (AFDC) was one of the largest federal assistance programs in the United States between 1935 and 1996. It consisted of a means-tested income support scheme for low-income families with dependent children, administered at the state level, but funded at the federal level. Following criticism that this program discouraged labor market participation and perpetuated welfare dependency, the Clinton administration enacted the 1996 Personal Responsibility and Work Opportunity Reconciliation Act (PRWORA), requiring all US states to replace AFDC with a Temporary Assistance for Needy Families (TANF) program. TANF programs differed amongst US states and were all fundamentally different from AFDC: they included strict time limits for the receipt of benefits and, simultaneously, generous earnings disregard schemes to incentivize work.

Under the policy framework of TANF, the state of Connecticut launched its own program, called Jobs First (JF) in 1996: this included the strictest time limit and also the most generous earnings disregard of all US states. Nonetheless, there was a transition period during which a policy experiment was conducted by the Manpower Demonstration and Research Corporation (MDRC). A random sample of approximately 5000 welfare applicants was randomly assigned to one of two groups: half of them were assigned to JF and faced its eligibility and program rules; the other half were randomly assigned to AFDC (the program that JF aimed to replace in the state of Connecticut), thereby facing AFDC eligibility and program rules.

The MDRC experimental data include data on quarterly income for a pre-program assignment period and also for a post-program assignment period, thereby allowing one to quantify and compare the income gains and losses experienced by the households that were randomly assigned to JF and AFDC¹² [Bitler et al. \(2006\)](#) use these experimental data to compare the distribution of income between the beneficiaries of AFDC and JF. They find that, while JF made the majority of individuals better-off, it also made a significant number of them worse-off, especially after the JF time limit kicks in and becomes binding.¹³ In our simple empirical illustration we draw on [Bitler et al. \(2006\)](#) and consider “AFDC” and “JF” as our alternative policies (equivalent to policies *A* and *B* in the previous sections). We illustrate our methods by constructing a LASD partial order to support a policy choice between these two programs¹⁴ and comparing it with the partial ordering that would emerge if loss aversion were not taken into consideration using conventional first order stochastic dominance (FOSD).¹⁵ Along the lines of [Bitler et al. \(2006\)](#) we make this comparison separately for the time period up until the JF time limit for the receipt of welfare benefits becomes binding and for the period after that.

5.1. JF vs AFDC: LASD ordering

To make welfare decisions in terms of gains and losses, we use data on household income changes, i.e. the difference between households’ income after exposure to the program (JF or AFDC) and before exposure to that program. We make this analysis separately for the period before the JF time limit become binding (TL) and for after that. We thus call pre-TL observations those that were made after random assignment to either of the policies (JF or AFDC) but before the time limit; we call post-TL observations those made after the JF time limit. We summarize household income (for both policies and pre/post TL periods) by averaging

¹² [Bitler et al. \(2006\)](#) conduct a test comparing features of households before random assignment and find that they do not differ significantly in terms of observable characteristics. We check additionally that the income distributions were the same before the experiment split households among the two policies. We use a conventional two-sided Cramér-von Mises test for the equality of distributions. The statistic was approximately 0.78 and its *p*-value was 0.55, implying that before the experiment, the distributions are indistinguishable.

¹³ [Bitler et al. \(2006\)](#) focus on quantile treatment effects (QTEs). If QTEs were to be used as a measure of the impact on any individual household in a welfare comparison, it would require the assumption of rank invariance across potential outcome distributions, which would be quite strong. Note that [Bitler et al. \(2006\)](#) do not make this assumption.

¹⁴ Although not directly relevant for our empirical illustration, it can be mentioned that the debate on the replacement of AFDC by TANF combined political economy concerns and also normative considerations about the appropriateness of policy-makers causing income losses to parts of the population. [Alesina et al. \(2001\)](#) use the AFDC as an empirical proxy for the generosity of the welfare state in the US and show that changes to this program had the potential to sway the electorate. At the same time, normative arguments supporting policy-makers’ loss-aversion have also been put forth in this context. Peter Edelman, then a senior advisor to President Clinton, resigned in protest against this policy change, calling the replacement of AFDC by TANF a “crucial moral litmus test”, as it risked causing important income losses to some households.

¹⁵ Because assignment is random, we assume that the distribution functions of gains and losses under each policy, F_{JF} and F_{AFDC} , are point-identified by the differences in incomes before and after random assignment.

Table 1

Tests for inferring whether the Jobs First (JF) program would be preferred to the Aid to Families with Dependent Children (AFDC). Column titles paraphrase the null hypotheses in the tests. The first column uses changes in income and the second column measures income in levels without regard to pre-policy income. 1999 bootstrap repetitions used in each test.

	LASD	FOSD
	$F_{JF} \geq F_{AFDC}$	$G_{JF} \geq G_{AFDC}$
Before JF time limit	0.1790	0.2240
p-value	0.9095	0.8634
After JF time limit	9.1380	2.1285
p-value	0.0000	0.1436




Fig. 1. Empirical distribution functions of changes in income near the JF time limit and the $T_2(F)$ coordinate processes that are used to test loss aversion-sensitive dominance. The second and third panels show the two coordinate functions in $T_2(F_n)$ defined in Eq. (12). The large positive values in the second panel drive the rejection of the hypothesis $F_{JF} \geq F_{AFDC}$ seen in Table 1.

income over all quarters in the relevant time span.¹⁶ Changes in household income due to the AFDC and JF policies were defined as the natural logarithm of the average household income in all post-policy quarters (either JF or AFDC) minus the natural log of the average pre-policy quarterly household income. Thus, our analysis applies LASD to these changes in two separate periods, the pre-TL period and the post-TL one.

The left-hand side of Table 1 shows the results of formal tests of hypothesis (10) using W_{2n} statistics (Cramér-von Mises statistics applied to the empirical T_2 process).¹⁷ For the pre-time limit period we cannot reject the hypothesis that $F_{JF} \geq_{LASD} F_{AFDC}$.¹⁸ However, everything changes when we make this comparison taking into account the post-time limit period. As mentioned above, after the time limit becomes binding, Bitler et al. (2006) show that a sizeable number of households in JF experience total income losses, as they stop receiving welfare transfers; this does not happen amongst households on AFDC, which does not have a time limit. In order to rank the distribution of income changes under JF and AFDC using LASD we test the hypothesis that $F_{JF} \geq_{LASD} F_{AFDC}$. As shown in the left-hand side of Table 1, this hypothesis is rejected for every significance level, reflecting the greater weight placed on the income losses experienced by JF beneficiaries.

To investigate how this rejection occurs, Fig. 1 displays the CDFs of gains and losses under the AFDC and JF policies around the JF time limit, then the way that the two T_2 coordinate processes compare them — when looking at the coordinates in Eq. (12), large positive values correspond to a rejection of the hypothesis $F_{JF} \geq_{LASD} F_{AFDC}$. The positive parts of the m_1 and m_2 functions illustrated in the middle and right-hand plots of Fig. 1 are squared and integrated over estimated contact sets to arrive at the test statistic in the lower left of Table 1. It can be seen in the second and third panels that the presumable reason that the JF policy does not dominate the AFDC policy using LASD is because the distribution of small gains and losses is more appealing in the AFDC program and the relation between small gains and small losses is preferable to JF.

5.2. JF vs AFDC: first order stochastic dominance ordering

What difference would it make if loss-aversion had been left out of this welfare ordering of social policies? In order to address this question we compare the welfare ordering obtained in the previous section with that obtained by ordering JF and AFDC according

¹⁶ We explored alternative definitions of our outcome of interest such as using the final quarter within the time span; generally these led to the same results, so we will not show them for the purpose of this simple illustration.

¹⁷ Results for the other test statistics are qualitatively the same. They are collected in an Online Supplemental Appendix.

¹⁸ For this time period we cannot even reject the null of equality in the distributions of changes in income between households assigned to JF and AFDC.




Fig. 2. Empirical distribution functions of (levels of) post-TL income and the resulting process that is used to test first-order stochastic dominance. Although there are positive values in the second panel, they are not large enough to drive a rejection of the hypothesis $G_{JF} \succeq_{FOSD} G_{AFDC}$, as seen in Table 1.

to first order stochastic dominance (FOSD). Using FOSD, the only relevant comparison is between the post-policy household income under JF and AFDC (household income before exposure to these policies is not material). We thus define our outcome of interest in levels, i.e. the natural log of the average household income under JF or AFDC. As before, we do this analysis separately for the two relevant time periods: before the JF time limit becomes binding and after it does.

The right-hand side of Table 1 shows the result of our FOSD tests. Either way post-policy outcomes are measured, we cannot reject the null that $G_{JF} \succeq_{FOSD} G_{AFDC}$. When measurements are made before and after exposure to the policy this is unsurprising, as prior to the JF time limit becoming binding none of the policies produces large income losses. However, even after the JF time limit becomes binding, the FOSD test still does not allow us to reject $G_{JF} \succeq_{FOSD} G_{AFDC}$, while the LASD test would lead us to categorically reject the dominance of JF over AFDC. This simple empirical illustration shows that, in practice, the consideration of loss aversion can change the welfare ordering of social policies.

Fig. 2 shows an analogous investigation into the way the analysis would typically be conducted using first order stochastic dominance to compare outcomes, using contact sets as in Linton et al. (2010). The contact set was estimated using $a_n = 4 \log(\log(n))$, corresponding to the tuning parameter choice of that paper, for both the FOSD and LASD tests (the smaller sequence $c_n = \sqrt{\log(\log(n))}$ was used for estimating near-maximizing sets in LASD tests). The left-hand plot in the figure shows the empirical distribution functions of outcomes under each program after the JF time limit. That is, the functions are based on levels of income rather than changes in income. The scaled difference $\sqrt{n}(G_{n,JF} - G_{n,AFDC})$ is displayed as the criterion function on the right-hand side of the panel, and the square of the positive part of this function is integrated over an estimated contact set. As the lower-right test in Table 1 indicates, these differences are sometimes mildly positive, so that G_{AFDC} is occasionally below G_{JF} (especially at lower income levels) but the difference is not large enough to indicate a rejection of the hypothesis that $G_{JF} \succeq_{FOSD} G_{AFDC}$, as indicated by the p -value of the test. Because outcomes are measured in levels, there is no way to measure whether they represent gains or losses for agents, and so a simple difference is used here instead of the comparison that accounts for loss aversion used with changes in income.

6. Conclusion

Public policies often result in gains for some individuals and losses for others. We define a social preference relation for distributions of gains and losses caused by a policy: loss aversion-sensitive dominance (LASD). We relate these social preferences to criteria that depend solely on distribution functions. The assumption of loss aversion can lead to a welfare ranking of policies that is different from the one that would be brought about if classic utility theory and first-order stochastic dominance were used. We then propose empirically testable conditions for LASD based on our CDF-based criterion functions. Because data may come as differences between underlying random variables, we propose a point-identified version of these conditions and also a partially identified analog. We develop inference methods to formally test LASD relations and derive the corresponding statistical properties. We show that resampling techniques, tailored to specific features of the criterion functions, can be used to conduct inference. Finally, we illustrate our LASD criterion and inference methods with a simple empirical application that uses data from a well known evaluation of a large income support policy in the US. This shows that the ranking of policy options depends on whether or not one takes individual loss aversion into account.

Appendix A. Proof of results

A.1. Results in Section 2

Proof of Proposition 2.3. Eq. (1) implies that

$$W(F) = \int_{\mathbb{R}_-} v(x) dF(x) + \int_{\mathbb{R}_+} v(x) dF(x). \quad (34)$$

We will now re-write the two parts of (34) using integration by parts, normalization $v(0) = 0$ and the fact that F has bounded support i.e. there exist $R_1 > 0$ and $R_2 < 0$ such that $\int_{R_2}^{R_1} f(x) dx = 1$, in which case we have $1 - F(R_2) = 0$ and $F(R_1) = 0$. For the first part of (34) we obtain

$$\begin{aligned} \int_{\mathbb{R}_-} v(x) dF(x) &= \lim_{R \rightarrow -\infty} \int_R^0 v(x) dF(x) \\ &= \lim_{R \rightarrow -\infty} \left[v(x)F(x) \Big|_R^0 - \int_R^0 v'(x)F(x) dx \right] \\ &= - \int_{-\infty}^0 v'(x)F(x) dx \end{aligned}$$

and for the second part (34) we have

$$\begin{aligned} \int_{\mathbb{R}_+} v(x) dF(x) &= - \int_{\mathbb{R}_+} v(x) d(1 - F)(x) \\ &= - \lim_{R \rightarrow \infty} \int_0^R v(x) d(1 - F)(x) \\ &= - \lim_{R \rightarrow \infty} \left[v(x)(1 - F(x)) \Big|_0^R - \int_0^R v'(x)(1 - F(x)) dx \right] \\ &= \int_0^\infty v'(x)(1 - F(x)) dx. \end{aligned}$$

Putting these two parts together yields (2). $\square$

A.2. Proofs of results in Section 3

Proof of Theorem 3.1. Notice that (3) is equivalent to both (4) and (5); in this proof we use the latter two conditions. Using Proposition 2.3 we rewrite $W(F_A) \geq W(F_B)$ as the equivalent condition

$$- \int_{-\infty}^0 v'(z)F_A(z) dz + \int_0^\infty v'(z)(1 - F_A(z)) dz \geq - \int_{-\infty}^0 v'(z)F_B(z) dz + \int_0^\infty v'(z)(1 - F_B(z)) dz.$$

Rearranging terms we find this is equivalent to

$$\int_{-\infty}^0 v'(z)F_B(z) dz - \int_{-\infty}^0 v'(z)F_A(z) dz \geq \int_0^\infty v'(z)(1 - F_B(z)) dz - \int_0^\infty v'(z)(1 - F_A(z)) dz$$

or simply

$$\int_{-\infty}^0 v'(z)(F_B(z) - F_A(z)) dz \geq \int_0^\infty v'(z)(F_A(z) - F_B(z)) dz.$$

This is in turn equivalent to

$$\int_0^\infty v'(-z)(F_B(-z) - F_A(-z)) dz \geq \int_0^\infty v'(z)(F_A(z) - F_B(z)) dz$$

or

$$- \int_0^\infty v'(-z)(F_A(-z) - F_B(-z)) dz \geq \int_0^\infty v'(z)(F_A(z) - F_B(z)) dz.$$

Adding $v'(z)(F_A(-z) - F_B(-z))$ to both sides we find this is equivalent to

$$\int_0^\infty (v'(z) - v'(-z))(F_A(-z) - F_B(-z)) dz \geq \int_0^\infty v'(z)(F_A(z) - F_B(z) + F_A(-z) - F_B(-z)) dz. \quad (35)$$

Utilizing the assumptions of loss aversion and non-decreasingness given in Definition 2.2, (4) and (5) are sufficient for (35) to hold for any v . Condition (5) is due to the fact that

$$F_A(x) - F_B(x) + F_A(-x) - F_B(-x) \leq 0 \quad \forall x \geq 0$$

is equivalent to the condition

$$1 - F_A(x) - F_A(-x) \geq 1 - F_B(x) - F_B(-x) \quad \forall x \geq 0.$$

We now show that conditions (4) and (5) are also necessary by means of a contradiction to (35). To this end, assume that there exists some $x > 0$ such that $F_A(-x) - F_B(-x) > 0$. From the fact that the distribution function is right continuous, it follows that there is a neighborhood (a, b) , $b > a > 0$, such that for all $x \in (a, b)$, $F_A(-x) - F_B(-x) > 0$. For arbitrarily small $\epsilon > 0$, consider the value function

$$v_1(x) = \begin{cases} a - b & x \leq -b - \epsilon \\ \frac{1}{4\epsilon}(x + b + \epsilon)^2 + a - b & x \in (-b - \epsilon, -b + \epsilon) \\ x + a & x \in (-b + \epsilon, -a - \epsilon) \\ \frac{1}{4\epsilon}(x + a - \epsilon)^2 & x \in (-a - \epsilon, -a + \epsilon) \\ 0 & x \geq -a + \epsilon. \end{cases}$$

Note that this v_1 satisfies Definition 2.2. Further, for $x \in (a, b)$, $v'_1(-x) > 0 = v'_1(x)$. Therefore

$$\int_0^\infty (v'_1(z) - v'_1(-z))(F_A(-z) - F_B(-z))dz < 0,$$

while

$$\int_0^\infty v'_1(z)(F_A(z) - F_B(z) + F_A(-z) - F_B(-z))dz = 0,$$

because $v'_1(x) = 0$ for $x \geq 0$. This contradicts (35).

The second condition can be proven similarly. Assume that there exists a neighborhood (a, b) , $0 < a < b$ such that for all $x \in (a, b)$, $(1 - F_A(x)) - F_A(-x) < (1 - F_B(x)) - F_B(-x)$. Take $v_2(x) = \text{sgn}(x) \times v_1(x)$. Using v_2 we find

$$\int_0^\infty (v'_2(z) - v'_2(-z))(F_A(-z) - F_B(-z))dz = 0$$

while

$$\int_0^\infty v'_2(z)(F_A(z) - F_B(z) + F_A(-z) - F_B(-z))dz > 0,$$

which is a contradiction. $\square$

Proof of Corollary 3.5. We first notice that $F_A \succeq_{FOSD} F_{SQ}$ is equivalent to the event

$$\{F_A \text{ is supported on } \mathbb{R}_+\}. \quad (36)$$

Property (36) easily implies that $F_A \succeq_{LASD} F_{SQ}$, which follows by Property 1 of Definition 2.2. On the other hand one checks that

$$v(x) := \begin{cases} x & x \leq 0 \\ 0 & x > 0 \end{cases}$$

fulfills Definition 2.2. Thus $F_A \succeq_{LASD} F_{SQ}$ implies (36). $\square$

Proof of Theorem 3.6. Given the bounds inequality, we have

$$L_B(-x) - U_A(-x) \leq F_B(-x) - F_A(-x) \leq U_B(-x) - L_A(-x)$$

and

$$L_A(x) - U_B(x) \leq F_A(x) - F_B(x) \leq U_A(x) - L_B(x),$$

from which it is clear that (8) is a sufficient condition. As a necessary condition we have (9) because using (3) we have

$$U_B(-x) - L_A(-x) \geq F_B(-x) - F_A(-x) \geq \max\{0, F_A(x) - F_B(x)\} \geq \max\{0, L_A(x) - L_B(x)\}. \quad \square$$

Proof of Corollary 3.7. Recall Corollary 3.5 implied that when F_B is a status quo distribution, the FOSD and LASD relations are equivalent. Then $F_A \succeq_{FOSD} F_{SQ}$ implies that $F_A(-x) = 0$ for all $x \geq 0$ because $F_{SQ}(-x) = 0$ for all $x \geq 0$. Therefore a sufficient condition for $F_A \succeq_{LASD} F_{SQ}$ is that $U_A(-x) = 0$ for all $x \geq 0$. Similarly, if $F_A \succeq_{LASD} F_{SQ}$, equivalent to $F_A \succeq_{FOSD} F_{SQ}$, then it must be the case that $F_A(-x) = 0$ for all $x \geq 0$, implying that $L_A(-x) = 0$ as well. $\square$

A.3. Results in Section 4

Proof of Theorem 4.1. For Part 1 note that if $\mathcal{P} \in \mathcal{P}_0$ then by definition, $\mathcal{X}_0^k(\mathcal{P}) \neq \emptyset$ for some $k \in \{1, 2\}$ and for all $x \in \mathcal{X}_0^k(\mathcal{P})$, $m_k(x) = 0$. For any k such that $\mathcal{X}_0^k(\mathcal{P}) \neq \emptyset$, that set is the limit of $\mathcal{M}_k(\epsilon)$ in the Painlevé-Kuratowski sense as $\epsilon \searrow 0$, because $\mathcal{M}_k(\epsilon)$ form a monotone sequence of sets and Exercise 4.3 of Rockafellar and Wets (1998) implies the limit is $\mathcal{X}_0^k(\mathcal{P})$. Then the supremum is achieved and $\lim_{\epsilon \searrow 0} \mathcal{M}_k(\epsilon) = \mathcal{X}_0^k(\mathcal{P})$ (in the Painlevé-Kuratowski sense) for at least one coordinate, so that suprema are taken over at least one of $\mathcal{X}_0^1(\mathcal{P})$ and $\mathcal{X}_0^2(\mathcal{P})$ and whichever coordinate satisfies this condition will contribute to the asymptotic distribution.

Note that for all $x \in \mathcal{X}_0(P)$, $\sqrt{n}T_1(\mathbb{F}_n)(x) = \sqrt{n}(T_1(\mathbb{F}_n) - T_1(F))(x)$. Lemma D.3 and the null hypothesis, which implies $\mathcal{X}_0^k(P) \neq \emptyset$ for $k \in \{1, 2\}$, imply the result for V_1 and W_1 .

To show Part 2, note that T_2 is a linear map of F , and assuming that $\mathcal{X}_0^k(P) \neq \emptyset$ for $k \in \{1, 2\}$, we have that its weak limit (for whichever set is nonempty) is $\sup_{x \in \mathcal{X}_0^k(P)} (T_{2k}(G_F)(x))^+$ by Lemma D.3. Breaking $\mathcal{X}_0(P)$ into its two subsets and assuming the null hypothesis is true results in the same behavior as the supremum norm statistic from the first part (using the definition of the supremum norm in two coordinates as the maximum of the two suprema). The same reasoning holds for the L_2 statistic in Part 2.

Part 3 follows from the behavior of the test statistics over $\{x \in \mathcal{X} : m_1(x) < 0, m_2(x) < 0\}$ described in Lemma D.3. To show Part 4 for V_{1n} suppose that for some x^* , $T_1(F)(x^*) = \xi > 0$. Then $\sup_{x \in \mathcal{X}} \sqrt{n}T_1(\mathbb{F}_n)(x) \geq \sqrt{n}(T_1(\mathbb{F}_n)(x^*) - T_1(F)(x^*)) + \sqrt{n}\xi$. Then

$$\begin{aligned} \liminf_{n \rightarrow \infty} P \left\{ \sup_{x \geq 0} \sqrt{n}T_1(\mathbb{F}_n)(x) > c \right\} \\ \geq \lim_{n \rightarrow \infty} P \left\{ \sqrt{n}(T_1(\mathbb{F}_n)(x^*) - T_1(F)(x^*)) > c - \sqrt{n}\xi \right\} \rightarrow 1, \end{aligned}$$

where the last convergence follows from the delta method applied to $\sqrt{n}(\mathbb{F}_n(x^*) - F(x^*))$, which converges in distribution to a tight random variable. The proof for the other statistics is analogous. $\square$

Proof of Theorem 4.2. This theorem is an application of Theorems 3.2 and 3.3 of Fang and Santos (2019). Define the statistics V_1 and W_1 as maps from F to the real line using ν and ω defined in equations (55) and (56) in Lemma D.3, and let their estimators be defined as in part 3 of the resampling scheme. Their Assumptions 1–3 are satisfied either by the definitions of ν and ω and Lemma D.3, the standard convergence result $\sqrt{n}(\mathbb{F}_n - F) \rightsquigarrow \mathcal{G}_F$ (van der Vaart and Wellner, 1996, Theorem 2.8.4) and the choice of bootstrap weights. We need to show that their Assumption 4 is also satisfied. Write either function as $\|h_1^+\| + \|h_1^+ \vee h_2^+\| + \|h_2^+\|$ using the desired norm. Both norms satisfy a reverse triangle inequality, and using the fact that $|(x)^+ - (y)^+| \leq |x - y|$, the difference for two functions g and h is bounded by $\|g_1 - h_1\| + \|g_1 \vee g_2 - h_1 \vee h_2\| + \|g_2 - h_2\|$. The first difference is bounded by $2\|g - h\|$, and the second and the third are bounded by $4\|g - h\|$. Rewriting Eqs. (31) and (32) as functionals of differential directions h , define

$$\hat{\nu}'_n(h) = \begin{cases} \left(\max_{x \in \hat{\mathcal{X}}_1^k} h_k(x) \right)^+ & |\max \hat{m}_{1n} - \max \hat{m}_{2n}| > c_n \\ \max \{0, \max_{x \in \hat{\mathcal{X}}_1^1} h_1(x), \max_{x \in \hat{\mathcal{X}}_2^2} h_2(x)\} & |\max \hat{m}_{1n} - \max \hat{m}_{2n}| \leq c_n \end{cases}$$

and

$$\hat{\omega}'_n(h) = \left(\int_{\hat{\mathcal{X}}_0^1} \left((h_1(x))^+ \right)^2 dx + \int_{\hat{\mathcal{X}}_0^2} \left((h_2(x))^+ \right)^2 dx \right)^{1/2}. \quad (37)$$

Because both ν and ω are Lipschitz, Lemma S.3.6 of Fang and Santos (2019) implies we need only check that $|\hat{\nu}'_n(h) - \nu'_F(h)| = o_P(1)$ and $|\hat{\omega}'_n(h) - \omega'_F(h)| = o_P(1)$ for each fixed h . This follows from the consistency of the contact set and ϵ -argmax estimators. The consistency of these estimators follow from the uniform law of large numbers for the ϵ -maximizing sets, and the tightness of the limit $\mathcal{G}_F$ for the contact sets, which implies that $\lim_n P \left\{ \sqrt{n}\|\mathbb{F}_n - F\|_\infty \leq \sqrt{na_n} \right\} = 1$. $\square$

Appendix B. Supplementary data

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.jeconom.2023.105643>.

References

- Aaberge, R., Havnes, T., Mogstad, M., 2021. Ranking intersecting distribution functions. *J. Appl. Econometrics* 36, 639–662.
- Alesina, A., Glaeser, E., Sacerdote, B., 2001. Why doesn't the US have a European-style welfare system. In: NBER working paper 8524.
- Alesina, A., Passarelli, F., 2019. Loss aversion, politics and redistribution. *Am. J. Polit. Sci.* 63, 936–947.
- Andrews, D.W., Shi, X., 2013. Inference based on conditional moment inequalities. *Econometrica* 81, 609–666.
- Atkinson, A.B., 1970. On the measurement of inequality. *J. Econom. Theory* 2, 244–263.
- Barrett, G.F., Donald, S.G., 2003. Consistent tests for stochastic dominance. *Econometrica* 71, 71–104.
- Bhattacharya, D., Dupas, P., 2012. Inferring welfare maximizing treatment assignment under budget constraints. *J. Econometrics* 167, 168–196.
- Bitler, M.P., Gelbach, J.B., Hoynes, H.W., 2006. What mean impacts miss: Distributional effects of welfare reform experiments. *Amer. Econ. Rev.* 96, 988–1012.
- Carneiro, P., Hansen, K.T., Heckman, J.J., 2001. Removing the veil of ignorance in assessing the distributional impacts of social policies. *Swed. Econ. Policy Rev.* 8, 273–301.
- Cattaneo, M.D., Jansson, M., Nagasawa, K., 2020. Bootstrap-based inference for cube root asymptotics. *Econometrica* 88, 2203–2219.
- Chetverikov, D., Santos, A., Shaikh, A.M., 2018. The econometrics of shape restrictions. *Annu. Rev. Econ.* 10, 31–63.
- Chew, S.H., 1983. A generalization of the quasilinear mean with applications to the measurement of income inequality and decision theory resolving the Allais paradox. *Econometrica* 51, 1065–1092.
- Cho, J.S., White, H., 2018. Directionally differentiable econometric models. *Econom. Theory* 34, 1101–1131.
- Christensen, T., Connault, B., 2019. Counterfactual sensitivity and robustness. In: Working paper.
- Dehejia, R., 2005. Program evaluation as a decision problem. *J. Econometrics* 125, 141–173.
- Easterly, W., 2009. The burden of proof should be on interventionists – Doubt is a superb reason for inaction. *Boston Review – series Development in dangerous places*.
- Eckhoudt, L., Schlesinger, H., 2006. Putting risk in its proper place. *Amer. Econ. Rev.* 96, 280–289.
- Eyal, N., 2020. Why challenge trials of SARS-CoV-2 vaccines could be ethical despite risk of severe adverse events. *Ethics & Human Research*.

- Fan, Y., Guerre, E., Zhu, D., 2017. Partial identification of functionals of the joint distribution of “Potential Outcomes”. *J. Econometrics* 197, 42–59.
- Fan, Y., Park, S.S., 2010. Sharp bounds on the distribution of treatment effects and their statistical inference. *Econom. Theory* 26, 931–951.
- Fan, Y., Park, S.S., 2012. Confidence intervals for the quantile of treatment effects in randomized experiments. *J. Econometrics* 167, 330–344.
- Fang, Z., Santos, A., 2019. Inference on directionally differentiable functions. *Rev. Econom. Stud.* 86, 377–412.
- Firpo, S., Galvao, A.F., Parker, T., 2023. Uniform inference for value functions. *J. Econometrics* 235, 1680–1699.
- Firpo, S., Ridder, G., 2008. Bounds on functionals of the distribution of treatment effects. In: IEPW Working Paper No. 08-09.
- Firpo, S., Ridder, G., 2019. Partial identification of the treatment effect distribution and its functionals. *J. Econometrics* 213, 210–234.
- Fishburn, P.C., 1980. Continua of stochastic dominance relations for unbounded probability distributions. *J. Math. Econom.* 7, 271–285.
- Frank, M.J., Nelsen, R.B., Schweizer, B., 1987. Best-possible bounds for the distribution of a sum — A problem of Kolmogorov. *Probab. Theory Related Fields* 74, 199–211.
- Freund, C., Özden, Ç., 2008. Trade policy and loss aversion. *Amer. Econ. Rev.* 98, 1675–1691.
- Gajdos, T., Weymark, J.A., 2012. Introduction to inequality and risk. *J. Econom. Theory* 147, 1313–1330.
- Heckman, J.J., Smith, J.A., 1998. Evaluating the welfare state. In: Strom, S. (Ed.), *Econometrics and Economic Theory in the Twentieth Century: The Ragnar Frisch Centennial Symposium*. Cambridge University Press, New York.
- Heckman, J.J., Smith, J., Clements, N., 1997. Making the most out of programme evaluations and social experiments: Accounting for heterogeneity in programme impacts. *Rev. Econom. Stud.* 64, 487–535.
- Hirano, K., Porter, J., 2009. Asymptotics for statistical treatment rules. *Econometrica* 77, 1683–1701.
- Hong, H., Li, J., 2018. The Numerical Delta method. *J. Econometrics* 206, 379–394.
- Kahneman, D., Tversky, A., 1979. Prospect theory: An analysis of decision under risk. *Econometrica* 47, 263–292.
- Kasy, M., 2016. Partial identification, distributional preferences, and the Welfare ranking of policies. *Rev. Econ. Stat.* 98, 111–131.
- Kitagawa, T., Tetenov, A., 2018. Who should be treated? Empirical welfare maximization methods for treatment choice. *Econometrica* 86, 591–616.
- Kitagawa, T., Tetenov, A., 2021. Equality-minded treatment choice. *J. Bus. Econom. Statist.* 39, 561–574.
- Lee, S., Song, K., Whang, Y.-J., 2018. Testing for a general class of functional inequalities. *Econom. Theory* 34, 1018–1064.
- Levy, H., 2016. *Stochastic Dominance: Investment Decision Making under Uncertainty*, 3rd Edition. Springer International Publishing, Switzerland.
- Linton, O., Maasoumi, E., Whang, Y.-J., 2005. Consistent testing for stochastic dominance under general sampling schemes. *Rev. Econom. Stud.* 72, 735–765.
- Linton, O., Song, K., Whang, Y.-J., 2010. An improved bootstrap test of stochastic dominance. *J. Econometrics* 154, 186–202.
- Makarov, G., 1982. Estimates for the distribution function of a sum of two random variables when the marginal distributions are fixed. *Theory Probab. Appl.* 26 (4), 803–806.
- Mankiw, G., 2014. When the scientist is also a philosopher. *New York Times – Business*.
- Manski, C.F., 2004. Statistical treatment rules for heterogeneous populations. *Econometrica* 72, 1221–1246.
- Masten, M.A., Poirier, A., 2020. Inference on breakdown frontiers. *Quant. Econ.* 11, 41–111.
- Rabin, M., Thaler, R.H., 2001. Anomalies: Risk aversion. *J. Econ. Perspect.* 15, 219–232.
- Rick, S., 2011. Losses, gains, and brains: Neuroeconomics can help to answer open questions about loss aversion. *J. Consumer Psychol.* 21, 453–463.
- Rockafellar, R.T., Wets, R.J.-B., 1998. *Variational Analysis*. Springer-Verlag, Berlin.
- Roemer, J.E., 1998. *Theories of Distributive Justice*. Harvard University Press, Cambridge.
- Rüschendorf, L., 1982. Random variables with maximum sums. *Adv. Appl. Probab.* 14, 623–632.
- Samuelson, W., Zeckhauser, R., 1988. Status quo bias in decision making. *J. Risk Uncertain.* 1, 7–59.
- Sen, A.K., 2000. *Freedom, Rationality and Social Choice: The Arrow Lectures and Other Essays*. Oxford University Press, Oxford.
- Shaked, M., Shanthikumar, G.J., 1994. *Stochastic Orders and Their Applications*. Academic Press, San Diego, CA.
- Stoye, J., 2009. Minimax regret treatment choice with finite samples. *J. Econometrics* 151, 70–81.
- Tetenov, A., 2012. Statistical treatment choice based on asymmetric minimax regret criteria. *J. Econometrics* 166, 157–165.
- Tversky, A., Kahneman, D., 1991. Loss aversion in riskless choice: A reference-dependent model. *Q. J. Econ.* 106, 1039–1061.
- Tversky, A., Kahneman, D., 1992. Advances in prospect theory: Cumulative representation of uncertainty. *J. Risk Uncertain.* 5, 297–323.
- van der Vaart, A.W., 1998. *Asymptotic Statistics*. Cambridge University Press, Cambridge.
- van der Vaart, A.W., Wellner, J.A., 1996. *Weak Convergence and Empirical Processes*. Springer, New York.
- Wellner, J.A., 1992. Empirical processes in action: A review. *Internat. Statist. Rev.* 60, 247–269.
- Weymark, J.A., 1981. Generalized Gini inequality indices. *Math. Social Sci.* 1, 409–430.
- Williamson, R.C., Downs, T., 1990. Probabilistic arithmetic I. Numerical methods for calculating convolutions and dependency bounds. *Internat. J. Approx. Reason.* 4, 89–158.
- Yaari, M.E., 1987. The dual theory of choice under risk. *Econometrica* 55, 95–115.
- Yaari, M.E., 1988. A controversial proposal concerning inequality measurement. *J. Econom. Theory* 44, 381–397.