

INFERENCE ON SETS IN FINANCE

VICTOR CHERNOZHUKOV[†] EMRE KOCATULUM[§] KONRAD MENZEL[‡]

ABSTRACT. In this paper we consider the problem of inference on a class of sets describing a collection of admissible models as solutions to a single smooth inequality. Classical and recent examples include, among others, the Hansen-Jagannathan (HJ) sets of admissible stochastic discount factors, Markowitz-Fama (MF) sets of mean-variances for asset portfolio returns, and the set of structural elasticities in Chetty (2012)'s analysis of demand with optimization frictions. We show that the econometric structure of the problem allows us to construct convenient and powerful confidence regions based upon the weighted likelihood ratio and weighted Wald (directed weighted Hausdorff) statistics. The statistics we formulate differ (in part) from existing statistics in that they enforce either exact or first order equivariance to transformations of parameters, making them especially appealing in the target applications. Moreover, the resulting inference procedures are also more powerful than the structured projection methods, which rely upon building confidence sets for the frontier-determining sufficient parameters (e.g. frontier-spanning portfolios), and then projecting them to obtain confidence sets for HJ sets or MF sets. Lastly, the framework we put forward is also useful for analyzing intersection bounds, namely sets defined as solutions to multiple smooth inequalities, since multiple inequalities can be conservatively approximated by a single smooth inequality. We present two empirical examples that show how the new econometric methods are able to generate sharp economic conclusions.

Keywords: Hansen-Jagannathan bound, Markowitz-Fama bounds, Chetty bounds, Mean-Variance sets, Optimization Frictions, Inference, Confidence Set

JEL classification: C10, C50

Date: March 2008 - this version: November 20, 2012. This paper was first presented at the March 2008 CEMMAP conference on Inference in Partially Identified Models with Applications. We thank Enrique Sentana and Anna Mikusheva for pointing the relevance of structured projection approach as a benchmark for comparison, and Isaiah Andrews, Dan Greenwald, and Ye Luo for detailed comments on various drafts of the paper. We also thank Raj Chetty, Patrick Gagliardini, Raffaella Giacomini, Jerry Hausman, Olivier Scaillet, Nour Meddahi, Enno Mammen and participants of seminars at MIT, CEMMAP, 2009 Toulouse Conference on Financial Econometrics, and 2009 European Econometric Society Meetings in Milan for useful discussion and comments.

1. INTRODUCTION

In this paper we consider the problem of inference on a class of sets describing a collection of admissible models as solutions to a single smooth inequality:

$$\Theta_0 = \{\theta \in \Theta : m(\theta) \leq 0\}, \quad (1.1)$$

where $\Theta \subset \mathbb{R}^k$ is a compact parameter space, and $m : \mathbb{R}^k \mapsto \mathbb{R}$ is the inequality-generating smooth function, which is estimable from the data. This structure arises in a number of important examples in financial economics and public finance, as we explain below via a sequence of examples. We show that this structure leads to highly-tractable and powerful inference procedures, and demonstrate their usefulness in substantive empirical examples. Furthermore, the structure could be used to conservatively approximate more complicated problems, where the sets of interest are given as intersections of solutions to multiple smooth inequalities. (Indeed, in the latter case, we can conservatively approximate the multiple inequalities by a single smooth inequality. The benefits from doing so is the highly tractable, “regular” inference.)

Example 1: Mean-Variance Set for the Stochastic Discount Factor. In order to describe the problem, we recall Cochrane (2005)’s assertion that the science of asset pricing could be effectively summarized using the following two equations:

$$\begin{aligned} P_t &= E_t[M_{t+1}X_{t+1}], \\ M_{t+1} &= f(Z_{t+1}, \text{parameters}), \end{aligned}$$

where P_t is an asset price, X_{t+1} is the asset payoff, M_{t+1} is the stochastic discount factor (SDF) or pricing kernel (PK), which is a function f of some data Z_{t+1} and parameters, and E_t is the conditional expectation given information at time t . The set of SDFs M_t that can price existing assets generally form a proper set, i.e. a set that is not a singleton. SDFs are not unique as long as the existing payoffs to assets do not span the entire universe of possible random payoffs. Dynamic asset pricing models provide families of potential SDFs, for example, the standard consumption model predicts that an appropriate SDF can be stated in terms of the intertemporal marginal rate of substitution:

$$M_t = \beta \frac{u'(C_{t+1})}{u'(C_t)},$$

where u denotes a utility function parameterized by some parameters, C_t denotes consumption at time t , and β denotes the subjective discount factor. Note that the investor’s

optimal consumption plan, and therefore the marginal rate of substitution, generally depend on a set of additional state variables Z_t , including lifetime wealth and non-capital income. Hence, when markets are incomplete, asset prices alone do not in general pin down a unique stochastic discount factor across different aggregate states of the economy or across investors with different values of Z_t .

An important empirical problem in the context of the CAPM is to check which families of SDFs price the assets correctly and which do not. This reasoning forms the basis for many approaches to estimation and tests for particular specifications of an asset pricing model, most prominently Hansen and Singleton (1982)'s seminal analysis of the consumption-based CAPM, see also Ludvigson (2012) for a survey of recent developments on the subject. One leading approach for performing the check is to see whether the mean and standard deviation of SDFs (μ, σ) of the stochastic discount factor M_t are admissible. Let K be a compact convex body in \mathbb{R}^2 . The set of admissible means and standard deviations in K

$$\Theta_0 := \{ \text{admissible pairs } (\mu, \sigma^2) \in \mathbb{R}^2 \cap K \},$$

is introduced by Hansen and Jagannathan (1991) and known as the Hansen-Jagannathan (HJ) set. The boundary of that set $\partial\Theta_0$ is known as the HJ bound. In order to describe the admissible pairs in a canonical setting, let v and Σ denote the vector of mean returns and covariance matrix to assets $1, \dots, N$, which are assumed not to vary with information sets at each period t . In the following, we consider the “minimal sufficient parameters”:

$$S_{vv} = v'\Sigma^{-1}v, \quad S_{v1} = v'\Sigma^{-1}1_N, \quad S_{11} = 1'_N\Sigma^{-1}1_N, \quad (1.2)$$

where 1_N is a column vector of ones. Then the minimum variance $\sigma_{HJ}^2(\mu)$ achievable by a SDF given mean μ of the SDF is equal to

$$\sigma_{HJ}^2(\mu) = (1 - \mu v)'\Sigma^{-1}(1 - \mu v) = S_{vv}\mu^2 - 2S_{v1}\mu + S_{11}.$$

Therefore, the HJ set is equal to

$$\Theta_{HJ} = \{ \underbrace{(\mu, \sigma)}_{\theta} \in \underbrace{(\mathbb{R} \times \mathbb{R}_+) \cap K}_{\Theta} : \underbrace{\sigma_{HJ}(\mu) - \sigma}_{m(\theta)} \leq 0 \}.$$

That is,

$$\Theta_0 = \{ \theta \in \Theta : m_{HJ}(\theta) \leq 0 \}.$$

Note that the inequality-generating function $m_{HJ}(\theta)$ depends on the unknown means μ and covariance of returns Σ via the minimal sufficient parameters, $\gamma = (S_{vv}, S_{v1}, S_{11})'$, so that we will write $m_{HJ}(\theta) = m_{HJ}(\theta, \gamma)$.

Example 2: Mean-Variance Analysis of Asset Portfolios. Let us now describe the second problem. The classical Markowitz (1952) problem is to minimize the variance of a portfolio given some attainable level of return:

$$\min_w \mathbb{E}_t[r_{p,t+1} - \mathbb{E}_t[r_{p,t+1}]]^2 \quad \text{such that} \quad \mathbb{E}_t[r_{p,t+1}] = \mu,$$

where $r_{p,t+1}$ is the return of the portfolio, determined as $r_{p,t+1} = w'r_{t+1}$, and w is a vector of portfolio “weights” and r_{t+1} is a vector of returns on available assets. In a canonical version of the problem, we have that the vector of mean returns v and covariance of returns Σ do not vary with time period t , so that the problem becomes:

$$\sigma_M(\mu) = \min_w w'\Sigma w \quad \text{such that} \quad w'v = \mu.$$

An explicit solution for $\sigma_M(\mu)$ takes the form,

$$\sigma_M^2(\mu) = \frac{S_{11}\mu^2 - 2S_{v1}\mu + S_{vv}}{S_{vv}S_{11} - S_{v1}^2},$$

where the minimal sufficient parameters $\gamma = (S_{vv}, S_{v1}, S_{11})'$ are the same as in equation (1.2).

Therefore given a compact convex body $K \subset \mathbb{R}^2$, the Markowitz (M) set of admissible standard deviations and means in K is given by

$$\Theta_M = \{(\mu, \sigma) \in (\mathbb{R} \times \mathbb{R}_+) \cap K : \sigma_M(\mu) - \sigma \leq 0\} = \{\theta \in \Theta : m_M(\theta) \leq 0\},$$

where $m_M(\theta) := \sigma_M(\mu) - \sigma$. The boundary of the set Θ_M is known as the efficient frontier. Note that as in HJ example, the inequality-generating function $m_M(\theta) = m_M(\theta, \gamma)$ depends on the unknown parameters, the means and covariance of returns, $\gamma = (S_{vv}, S_{v1}, S_{11})'$. We also note that in some applications, the complementary Markowitz set is of interest:

$$\Theta_M^c = \{\theta \in \Theta : -m_M(\theta) \leq 0\}.$$

The inference on such sets thus also fall in our framework (1.1).

Mean-variance analysis for asset returns and the stochastic discount factor provides powerful tools to summarize the opportunities for risk diversification in a given asset market, and can serve as a basis for spanning and efficiency tests. The importance of

sampling error and the role of statistical inference in portfolio analysis has long been recognized in the literature: Gibbons et al. (1989) develop regression-based efficiency tests, and Britten-Jones (1999) considers inference on the tangency portfolio and tests for portfolio efficiency; recent work by Peñaranda and Sentana (2010) combines restrictions derived from the two dual approaches to mean-variance analysis in order to derive more powerful spanning tests for (sub-)markets; a summary of the recent literature on mean-variance efficiency tests is provided by Sentana (2009). The econometric contribution of this paper is quite distinct from and hence complementary to these previous efforts.

Example 3: Multi-Factor Efficient Frontiers. A third example due to Fama (1996) extends the unconditional portfolio choice problem in Markowitz (1952) to multiple sources of priced risk. More specifically, suppose that there exist k state variables $Z_t = (z_{1t}, \dots, z_{kt})'$ (“factors”) such that $(r'_t, Z'_t)'$ are jointly normal (or, more generally, jointly spherical). We assume that Z_t follows a K -variate standard normal (or other spherical) distribution, and

$$E_t[r_{t+1}|Z_{t+1}] = v + B'Z_{t+1}$$

where v is an N -dimensional vector of average returns, and the $k \times N$ covariance matrix $B = [\beta_1, \dots, \beta_N] = \text{Cov}_t(Z_{t+1}, r_{t+1})$ gives the factor loadings for the vector of asset returns. As in the previous examples, we denote the (unconditional) mean of asset returns by $v \in \mathbb{R}^N$, and their covariance matrix with Σ . Note that for the purposes of this example, we take r_{it} to denote the excess return of asset i relative to a risk-free asset with return r_f .

This analysis is motivated by the investor’s lifetime consumption-investment problem, where the state variables Z_t denote other factors influencing prices for consumption goods and the household’s income from wage labor or entrepreneurial activity in period t . In particular the household not only faces a trade-off between mean and variance of portfolio returns, but may also wish to use her portfolio to insure against price or income shocks.

Specifically, we consider the problem constructing a minimum variance portfolio that targets a mean return μ^* and a vector β^* of loadings respect to the factors Z_t in the absence of borrowing constraints. Denoting $D := [v, B']$ and $\delta^* := (\mu^*, \beta^{*'})'$, the efficient portfolio weights w can be found by solving the program

$$\min_w w' \Sigma w \quad \text{s.t.} \quad D'w = \delta^*. \quad (1.3)$$

Note that since the investor has access to the risk-free asset and faces no short-sale constraints, the portfolio weights w are not required to be nonnegative or sum up to one. Solving this program yields efficient portfolio weights $w^* = \Sigma^{-1}D(D'\Sigma D)^{-1}\delta^*$ so that the lower bound on the variance of a portfolio with beta equal to β and mean μ is given by

$$\sigma_{MF}^2(\mu, \beta) := (w^*)'\Sigma w^* = \delta^{*'}(D'\Sigma D)^{-1}\delta^*.$$

The resulting multi-factor efficient frontier bounds the cone

$$\Theta_{MF} := \left\{ (\mu^*, \beta^{*'}, \sigma)' \in (\mathbb{R}^{k+1} \times \mathbb{R}_+) \cap K : \sigma_{MF}(\mu^*, \beta^*) - \sigma \leq 0 \right\}$$

of feasible mean-variance-beta combinations, where K is some compact convex body in \mathbb{R}^{k+2} . Again, this cone can also be characterized by an inequality condition on the moment function

$$m_{MF}(\theta, \gamma) := \sqrt{(\mu^*, \beta^{*'})'(D'\Sigma D)^{-1}(\mu^*, \beta^{*'})' - \sigma}$$

where $\theta = (\mu^*, \beta^{*'}, \sigma)'$ and $\gamma = \text{vec}(D'\Sigma D)$.

Extending the two-fund theorem, Fama (1996) shows that we can always construct $k+2$ different funds from the N assets that span the entire multi-factor efficient frontier. In the presence of heterogeneity in preferences, wealth and consumption possibilities, different investors will typically choose different points on the efficient cone, so that only the entire set constitutes an adequate representation of the opportunities to insure against different factor risks in the market.

Example 4: Bounds on Demand Elasticities. To illustrate the wider applicability of our methods, we also consider inference on demand elasticities in the presence of an optimization friction. Chetty (2012) considers a household's dynamic labor supply and consumption problem and proposes bounds on price elasticities that allow for optimization errors of unknown form but that are bounded in magnitude. His analysis assumes that preferences can be represented by quasilinear flow utilities and lead to Marshallian demand functions of the form

$$\log x_{it}^*(p) = \alpha - \varepsilon \log p + \nu_{it},$$

where p is the price of the good of interest, and ν_{it} a preference shock for household i in period t . Here ε is the (constant) "structural" demand elasticity corresponding to the first-best solution to the household's planning problem.

Let U_{it}^* be the discounted lifetime utility from following the optimal consumption plan and let $e_{it}(x)$ denote the minimal expenditure needed to finance a utility level U_{it}^* under the constraint that x_{it} is held fixed at x . The analysis then considers “small” optimization errors in the households’ observed choice of x_{it} such that for every period t , the average utility loss measured in terms of the difference in expenditure as a percentage of the optimal budget is bounded by some $\delta > 0$, i.e.

$$\frac{1}{N} \sum_{i=1}^N \frac{e_{it}(x_{it}) - e_{it}(x_{it}^*(p_t))}{p_t x_{it}^*} \leq \delta,$$

where the average is taken over a sample of households $i = 1, \dots, N$.

Chetty (2012) derives the lower and upper bounds ε_L and ε_U for the structural elasticities that are compatible an observed elasticity ε^o given a price change $\Delta \log p$, and in the presence of an optimization error of $\delta > 0$. This bound is given by the following two inequalities

$$\varepsilon^o = \varepsilon_U - 2 \frac{(2\varepsilon_U \delta)^{1/2}}{\Delta \log p} \quad \text{and} \quad \varepsilon^o = \varepsilon_L + 2 \frac{(2\varepsilon_L \delta)^{1/2}}{\Delta \log p}$$

where δ is assumed to be small.¹ With a few basic manipulations, we can obtain the following equivalent characterization of the (δ, ε) -set as a lower bound for the distortion needed to reconcile ε^o with a structural elasticity ε :

$$\delta \geq \frac{(\varepsilon - \varepsilon^o)^2 (\Delta \log p)^2}{8\varepsilon} =: \delta_{OF}(\varepsilon, \varepsilon^o, \Delta \log p) \quad \text{if and only if } \varepsilon \in [\varepsilon_L, \varepsilon_U]. \quad (1.4)$$

Chetty (2012) gives several interpretations for this optimization friction, which could result e.g. from adjustment costs or misperception of prices. In the context of the consumption-CAPM model, Lettau and Ludvigson (2009) conducted a similar sensitivity analysis, allowing for optimization errors in terms of the Euler equations characterizing the household’s optimal investment decision.

The resulting set of structural elasticity/optimization friction pairs (ε, δ) compatible with the observed elasticity ε^o is given by

$$\Theta_{OF} := \{(\varepsilon, \delta) \in \mathbb{R}_+^2 : \delta_{OF}(\varepsilon, \varepsilon^o, \Delta \log p) - \delta \leq 0\}.$$

¹See equations (12) and (13) on pages 983-985 in Chetty (2012)

The bounds in equation (1.4) can therefore also be expressed as an inequality condition on the moment function

$$m_{OF}(\theta, \gamma) := \frac{(\varepsilon - \varepsilon^o)^2 (\Delta \log p)^2}{8\varepsilon} - \delta,$$

where $\theta = (\varepsilon, \delta)$ and $\gamma = \varepsilon^o$, and $\Delta \log p$ will be regarded as a fixed design parameter. In applications, one might be interested in intersections bounds from several studies in order to obtain tighter bounds, as was done in Chetty (2012). As we discuss below and illustrate in the empirical example, it is possible to perform this intersection in a smooth fashion, retaining the general single-equation structure as given in (1.1).

1.1. Overview. The basic problem of this paper is to develop inference methods on sets defined by nonlinear inequality restrictions, as the HJ and M sets, while accounting for uncertainty in the estimation of parameters of the inequality-generating functions. Specifically, the problem is to construct a confidence region R such that

$$\lim_{n \rightarrow \infty} P\{\Theta_0 \subseteq R\} = 1 - \alpha, \quad d_H(\Theta_0, R) = O_p(1/\sqrt{n}),$$

where d_H is the Hausdorff distance, defined below in (2.17). We will construct confidence regions for the set Θ_0 using LR and Wald-type Statistics, building on and simultaneously enriching the approaches suggested in Chernozhukov et al. (2007) and in Beresteanu and Molinari (2008) and Molchanov (1998), respectively. We also would like to ensure that confidence regions R are as small as possible and converge to Θ_0 at the fastest attainable speed.

Once R is constructed, we can test any composite hypotheses involving the parameter θ without compromising the significance level. E.g. a typical application of the HJ bounds determines which combinations of Θ_e the first two moments of the SDF generated by a given family of economic models fall in the HJ set. Indeed, in that case $\Theta_e = \{\mu(\varrho), \sigma(\varrho), \varrho \in [0, \infty)\}$, where ϱ is the elasticity of power utility function, and we can check which values of γ give us overlap of Θ_e with the confidence set R for the HJ set.² Similar comments about applicability of our approach go through for the M and MF sets as well. We should also mention that our confidence regions R are also more powerful than the structured projection methods, which rely upon building confidence sets for the frontier-determining sufficient parameters (e.g. frontier-spanning portfolios), and then projecting them to obtain confidence sets for HJ sets or MF sets. We demonstrate this

²In cases where curve $\varrho \mapsto \mu(\varrho), \varrho \mapsto \sigma(\varrho)$ are estimated, we need to construct a confidence set R_e for Θ_e and then look whether (and where) R_e and R overlap to determine the plausible values of ϱ .

in the empirical section, where we compute the confidence sets for HJ sets based on our approach and based on structured projection approach, and show that our confidence interval is much smaller and lies strictly inside the confidence region based on structured projections. We provide a theoretical explanation to this phenomenon in Section 2.5

Our procedure for inference using weighted Wald-type statistics complements other approaches based on the directed Hausdorff distance suggested in Beresteanu and Molinari (2008) and Hausdorff distance in Molchanov (1998). By using weighting in the construction of the Wald-type statistics, we make this approach asymptotically invariant to parameter transformations, which results in noticeably sharper confidence sets, at least in the canonical empirical example that we will show. These invariance properties should be seen as complementary to questions of efficient estimation of the identified set as studied by Kaido and Santos (2011); indeed, many types of confidence sets will all be centered around the same efficient set estimate, and so we use invariance and precision considerations to select amongst various potential constructions of sets.

Critchley et al. (1996) analyzed the geometric structure of Wald-type tests with nonlinear restrictions for point-identified problems and pointed out the beneficial role of invariance of the distribution of the statistic with respect to reparameterizations. In regular problems, the effect of nonlinearities on inference is typically asymptotically negligible due to the delta method and its extensions. However, our results show that for inference on sets with a nontrivial diameter, the effect of nonlinearities is of the same order as that of sampling variation. In particular we propose inference procedures that rely on quantities that are asymptotically pivotal and (asymptotically) invariant to parameter transformation, and illustrate that procedures failing to meet these requirements lead to overly conservative inference. In particular, these conditions are not met by the classical Hausdorff distance except in very special cases.

Thus, our construction is of independent interest for this type of inference and is a useful complement to the work of Beresteanu and Molinari (2008) and Molchanov (1998). Furthermore, our results on formal validity of the bootstrap for the weighed LR-type and W-type statistics are also of independent interest (for the class of problems with structures that are similar to those studied here; see Beresteanu and Molinari (2008) and Kaido and Santos (2011) for related bootstrap results for Hausdorff statistics.)

The rest of the paper is organized as follows. In Section 2 we present our estimation and inference results. In Section 3 we present an empirical example, illustrating the constructions of confidence sets for HJ sets. In Section 4 we draw conclusions and

provide direction for further research. In the Appendix, we collect the proofs of the main results.

2. ESTIMATION AND INFERENCE RESULTS

2.1. Basic Constructions. We first introduce our basic framework. We have a real-valued inequality-generating function $m(\theta)$, and the set of interest is the solution of the inequalities generated by the function $m(\theta)$ over a compact parameter space $\Theta \subset \mathbb{R}^k$:

$$\Theta_0 = \{\theta \in \Theta : m(\theta) \leq 0\}.$$

A natural estimator of Θ_0 is its empirical analog

$$\widehat{\Theta}_0 = \{\theta \in \Theta : \widehat{m}(\theta) \leq 0\},$$

where $\widehat{m}(\theta)$ is the estimate of the inequality-generating function. For example, in the HJ and M examples, the estimate takes the form

$$\widehat{m}(\theta) = m(\theta, \widehat{\gamma}), \quad \widehat{\gamma} = (\widehat{S}_{vv}, \widehat{S}_{v1}, \widehat{S}_{11})'.$$

Remark 1 (Approximating Multiple Inequalities by A Single Smooth Inequality). Throughout the paper, we will only consider the case of a single inequality restriction, $m : \Theta \rightarrow \mathbb{R}$ satisfying certain smoothness conditions with respect to θ . However this framework approximately encompasses the case of multiple inequalities, since it is possible to conservatively approximate intersection bounds of the form $g(\theta) \leq 0$, where $g(\theta) := (g_1(\theta), \dots, g_J(\theta))'$ represents a vector of constraints, by a single smooth bound $m(\theta) \leq 0$. To be specific, we can form an inequality-generating function:

$$m_\lambda(\theta) := m(g_1(\theta), \dots, g_J(\theta); \lambda) := \sum_{j=1}^J \frac{\exp(\lambda g_j(\theta))}{\sum_{l=1}^J \exp(\lambda g_l(\theta))} g_j(\theta),$$

where $\lambda \in \mathbb{R}$ is a fixed, positive scalar. This function m_λ is clearly smooth with respect to the underlying function g , and is a conservative approximation to $\max_{1 \leq j \leq J} g_j(\theta)$ with an explicit error bound, as shown by the following lemma.

Lemma 1 (Properties of Smooth Max). *(i) We have that*

$$m(g_1, \dots, g_J, \lambda) \leq 0 \quad \text{if} \quad \max_{1 \leq j \leq J} g_j(\theta) \leq 0.$$

(ii) For any $\lambda > 0$, the approximation error obeys

$$\sup_{(g_1, \dots, g_J)' \in \mathbb{R}^J} |\max(g_1, \dots, g_J) - m(g_1, \dots, g_J, \lambda)| \leq \lambda^{-1} W \left(\frac{J-1}{e} \right)$$

where e is Euler's number, and the function $W : \mathbb{R} \mapsto \mathbb{R}$ is Lambert's product logarithm function, i.e. $W(z)$ is defined as the solution w of the equation $w \exp\{w\} = z$. In particular, $W(z) \geq 0$ if $z \geq 0$, and $W(z) \leq \log(z)$ if $z \geq e$.

This lemma implies that the approximation error relative to the max function (with respect to the sup-norm) is inversely proportional to the smoothing parameter λ , where the constant of proportionality grows only very slowly in J . In the paper we will state our results and assumptions on estimation and inference directly in terms of the scalar moment $m(\theta) = m_\lambda(\theta)$ and its empirical analog $\widehat{m}(\theta) := \widehat{m}_\lambda(\theta) := m(\widehat{g}_1(\theta), \dots, \widehat{g}_J(\theta); \lambda)$. We also consider the quantity $1/\lambda$ fixed at some small value. \square

Our proposals for confidence regions are based on (1) a LR-type statistic and (2) a Wald-type statistic. The LR-based confidence region is

$$R_{LR} := \left\{ \theta \in \Theta : \left[\sqrt{n} \widehat{m}(\theta) / \widehat{s}(\theta) \right]_+^2 \leq \widehat{k}(1 - \alpha) \right\}, \quad (2.1)$$

where $\widehat{s}(\theta)$ is a weighting function; ideally, the standard error of $\widehat{m}(\theta)$; and $\widehat{k}(1 - \alpha)$ is some suitable estimate of $k(1 - \alpha)$, the $(1 - \alpha)$ -quantile of the statistic

$$\mathcal{L}_n = \sup_{\theta \in \Theta_0} \left[\sqrt{n} \widehat{m}(\theta) / \widehat{s}(\theta) \right]_+^2. \quad (2.2)$$

Note that \mathcal{L}_n is a LR-type statistic, as in Chernozhukov et al. (2007).

Next we shall consider confidence regions based on inverting a Wald-type statistic. In order to eliminate boundary effects on the distribution of this statistic, we will need to assume that the moment conditions, and therefore the set estimator, are well-defined in a neighborhood of the parameter space. To this end, we define the distance of a point θ to a set $A \subset \mathbb{R}^k$ as

$$d(\theta, A) := \inf_{\theta' \in A} \|\theta - \theta'\|.$$

We let the set

$$\Theta^\delta := \{\theta \in \mathbb{R}^k : d(\theta, \Theta) \leq \delta\}$$

denote the δ -expansion of Θ in \mathbb{R}^k for any $\delta > 0$. We denote the natural set estimator for Θ_0 in the expansion of the parameter space by

$$\widehat{\Theta}_{0,\delta} := \{\theta \in \Theta^\delta : \widehat{m}(\theta) \leq 0\}.$$

We will see below that the behavior of the moment function near the boundary of the parameter space is in general relevant for the statistical behavior of our procedure. However, for the practical applications discussed in this paper, the restriction of the

parameter space Θ to a compact set is entirely for technical reasons, so that an extension of the definitions to Θ^δ is unproblematic.

Given these definitions, our Wald-based confidence region is

$$R_W := \{\theta \in \Theta : [\sqrt{nd}(\theta, \widehat{\Theta}_{0,\delta})/\widehat{w}(\theta)]^2 \leq \widehat{k}(1 - \alpha)\}, \quad (2.3)$$

where $\widehat{w}(\theta)$ is a weighting function, particular forms of which we will suggest later; and $\widehat{k}(1 - \alpha)$ is a suitable estimate of $k(1 - \alpha)$, the $(1 - \alpha)$ -quantile of \mathcal{W}_n , where \mathcal{W}_n is the weighted W-statistic

$$\mathcal{W}_n := \sup_{\theta \in \Theta_0} [\sqrt{nd}(\theta, \widehat{\Theta}_{0,\delta})/\widehat{w}(\theta)]^2. \quad (2.4)$$

In the special case, where the weight function is flat, namely $\widehat{w}(\theta) = w$ for all θ , the W-statistic \mathcal{W}_n becomes the canonical directed Hausdorff distance (see Molchanov (1998), and Beresteanu and Molinari (2008)):

$$\sqrt{\mathcal{W}_n} \propto d^+(\Theta_0, \widehat{\Theta}_{0,\delta}) = \sup_{\theta \in \Theta_0} \inf_{\theta' \in \widehat{\Theta}_{0,\delta}} \|\theta - \theta'\|.$$

Remark 2 (Invariance Motivation for Weighed W-Statistics). The weighted W-statistic (2.4) is generally not a distance, but we argue in Section 2.6 that it provides a very useful extension/generalization of the canonical directed Hausdorff distance. Note that, like any Euclidean norm, the Hausdorff distance is not invariant with respect to changes in the scale of the different components of θ , which results in confidence regions that are not equivariant to such transformations. Our empirical results below illustrate that as a result, the shape of confidence sets alters dramatically as we change the weight on different dimensions of the parameter space. In sharp contrast, the introduction of certain types of weights makes the W-statistic (first-order) invariant to such transformations, as shown in Section 2.6, which makes the resulting confidence regions (first-order) equivariant. As a result, in all of our empirical examples such weighting dramatically improves the confidence regions (e.g. compare Figures 5 and 6 in the example concerning inference on HJ sets.) \square

2.2. A Basic Limit Theorem for LR and W statistics. In this subsection, we develop a basic result on the limit laws of the LR and W statistics. We will develop this result under the following general regularity conditions:

CONDITION R. *For some $\delta > 0$, the inequality-generating functions $m : \Theta^\delta \rightarrow \mathbb{R}$ and its estimator $\widehat{m} : \Theta^\delta \rightarrow \mathbb{R}$ are well-defined.*

(R.1) *The estimator \widehat{m} is $n^{-1/2}$ -consistent for m and asymptotically Gaussian, namely, in the metric space of bounded functions $\ell^\infty(\Theta^\delta)$*

$$\sqrt{n}(\widehat{m} - m) \rightsquigarrow \mathbb{G},$$

where \mathbb{G} is a zero-mean Gaussian process with continuous paths a.s., and a non-degenerate covariance function, i.e. $\inf_{\theta \in \Theta^\delta} \mathbb{E}[\mathbb{G}^2(\theta)] > 0$.

(R.2) *Functions \widehat{m} and m admit continuous gradients $\nabla_\theta \widehat{m} : \Theta^\delta \mapsto \mathbb{R}^k$ and $\nabla_\theta m : \Theta^\delta \mapsto \mathbb{R}^k$ with probability one, where*

$$\nabla_\theta \widehat{m}(\theta) = \nabla_\theta m(\theta) + o_p(1).$$

uniformly in $\theta \in \Theta^\delta$. The gradient $\theta \mapsto \nabla_\theta m$ is uniformly Lipschitz and $\inf_{\theta \in \Theta^\delta} \|\nabla_\theta m(\theta)\| > 0$.

(R.3) *Weighting functions $\widehat{s} : \Theta^\delta \mapsto \mathbb{R}$ and $s : \Theta^\delta \mapsto \mathbb{R}$ satisfy uniformly in $\theta \in \Theta^\delta$*

$$\widehat{s}(\theta) = s(\theta) + o_p(1), \quad \widehat{w}(\theta) = w(\theta) + o_p(1),$$

where $s : \Theta^\delta \mapsto \mathbb{R}_+$ and $w : \Theta^\delta \mapsto \mathbb{R}_+$ are continuous functions with values bounded away from zero.

In Condition R.1, we require the estimates of the inequality-generating functions to satisfy a uniform central limit theorem. Many sufficient conditions for this are provided by the theory of empirical processes, see e.g. van der Vaart and Wellner (1996). In our finance examples, this condition will follow from asymptotic normality of the estimates of the mean returns and covariance of returns. In Condition R.2, we require that gradient of the estimate of the inequality-generating function be consistent for the gradient of the inequality-generating function. Moreover, we require that the norm $\|\nabla_\theta m(\theta)\|$ be bounded away from zero, which is an identification condition and allows us to estimate, at a parametric rate, the boundary of the set Θ_0 , which we define as

$$\partial\Theta_0 := \{\theta \in \Theta : m(\theta) = 0\}.$$

In Condition R.3, we require that the estimates of the weight functions be consistent, and the weight functions be well-behaved.

Under these conditions we can state the following general result.

Theorem 1. (Limit Laws of LR and W Statistics). *Under Condition R*

$$\mathcal{L}_n \rightsquigarrow \mathcal{L}, \quad \mathcal{L} = \sup_{\theta \in \partial\Theta_0} \left[\frac{\mathbb{G}(\theta)}{s(\theta)} \right]_+^2, \quad (2.5)$$

$$\mathcal{W}_n \rightsquigarrow \mathcal{W}, \quad \mathcal{W} = \sup_{\theta \in \partial\Theta_0} \left[\frac{\mathbb{G}(\theta)}{\|\nabla_{\theta} m(\theta)\| \cdot w(\theta)} \right]_+^2, \quad (2.6)$$

where both \mathcal{W} and \mathcal{L} have distribution functions that are continuous at their $(1 - \alpha)$ -quantiles for $\alpha < 1/2$. Furthermore, if

$$\widehat{w}(\theta) = \frac{s(\theta)}{\|\nabla_{\theta} m(\theta)\|} + o_p(1),$$

uniformly in $\theta \in \Theta^{\delta}$, then the two statistics are asymptotically equivalent:

$$\mathcal{W}_n = \mathcal{L}_n + o_p(1).$$

In particular, this equivalence occurs if $\widehat{w}(\theta) = \|\nabla_{\theta} \widehat{m}(\theta)\|/\widehat{s}(\theta)$.

We see from this theorem that the LR and W statistics converge in law to well-behaved random variables that are continuous transformations of the limit Gaussian process \mathbb{G} . Moreover, we see that under an appropriate choice of the weighting functions, the two statistics are asymptotically equivalent.

For our application to HJ and MF sets, the following conditions will be sufficient.

CONDITION C.

(C.1) *We have that $m(\theta) = m(\theta, \gamma_0)$, where $\gamma_0 \in \Gamma \subset \mathbb{R}^d$, for all $\theta \in \Theta^{\delta}$ and some $\delta > 0$. The value γ_0 is in the interior of Γ , and there is an estimator $\widehat{\gamma}$ of γ_0 that obeys*

$$\sqrt{n}(\widehat{\gamma} - \gamma_0) \rightsquigarrow \Omega^{1/2}Z, \quad Z = N(0, I_d).$$

for some positive-definite Hermitian matrix Ω . Moreover, there is $\widehat{\Omega} \rightarrow_p \Omega$.

(C.2) *The gradient map $(\theta, \gamma) \mapsto \nabla_{\theta} m(\theta, \gamma)$, mapping $\Theta^{\delta} \times \Gamma \mapsto \mathbb{R}^k$, exists and is uniformly Lipschitz-continuous. Moreover, $\inf_{(\theta, \gamma) \in \Theta^{\delta} \times \Gamma} \|\nabla_{\theta} m(\theta, \gamma)\| > 0$.*

(C.3) *The gradient map $(\theta, \gamma) \mapsto \nabla_{\gamma} m(\theta, \gamma)$, mapping $\Theta^{\delta} \times \Gamma \mapsto \mathbb{R}^k$, exists and is uniformly Lipschitz-continuous.*

We show in Proposition 1 that these conditions hold for the canonical versions of the HJ and MF problems. Under these conditions we immediately conclude that in the

metric space of bounded functions $\ell^\infty(\Theta^\delta)$:

$$\begin{aligned}\sqrt{n}(\widehat{m}(\cdot) - m(\cdot)) &= \nabla_\gamma m(\cdot, \bar{\gamma})' \sqrt{n}(\widehat{\gamma} - \gamma_0) + o_p(1) \\ &\rightsquigarrow \nabla_\gamma m(\cdot, \gamma_0)' \Omega^{1/2} Z,\end{aligned}$$

where $\nabla_\gamma m(\theta, \bar{\gamma})$ denotes the gradient with each of its rows evaluated at a value $\bar{\gamma}$ on the line connecting $\widehat{\gamma}$ and γ_0 , where value $\bar{\gamma}$ may vary from row to row of the matrix. Therefore, the limit process in HJ and M examples takes the form:

$$\mathbb{G}(\theta) = \nabla_\gamma m(\theta, \gamma_0)' \Omega^{1/2} Z. \quad (2.7)$$

This will lead us to conclude formally below that conclusions of Theorem 1 hold with

$$\mathcal{L} = \sup_{\theta \in \partial\Theta_0} \left[\frac{\nabla_\gamma m(\theta, \gamma_0)' \Omega^{1/2} Z}{s(\theta)} \right]_+^2, \quad (2.8)$$

$$\mathcal{W} = \sup_{\theta \in \partial\Theta_0} \left[\frac{\nabla_\gamma m(\theta, \gamma_0)' \Omega^{1/2} Z}{\|\nabla_\theta m(\theta, \gamma_0)\| \cdot w(\theta)} \right]_+^2. \quad (2.9)$$

A good strategy for choosing the weighting function for LR and W is to choose the studentizing Anderson-Darling weights

$$s(\theta) = \|\nabla_\gamma m(\theta, \gamma_0)' \Omega^{1/2}\|, \quad (2.10)$$

$$w(\theta) = \frac{\|\nabla_\gamma m(\theta, \gamma_0)' \Omega^{1/2}\|}{\|\nabla_\theta m(\theta, \gamma_0)\|}. \quad (2.11)$$

The natural estimates of these weighting functions are given by the following plug-in estimators:

$$\widehat{s}(\theta) := \|\nabla_\gamma m(\theta, \widehat{\gamma})' \widehat{\Omega}^{1/2}\|, \quad (2.12)$$

$$\widehat{w}(\theta) := \frac{\|\nabla_\gamma m(\theta, \widehat{\gamma})' \widehat{\Omega}^{1/2}\|}{\|\nabla_\theta m(\theta, \widehat{\gamma})\|}. \quad (2.13)$$

We formalize the preceding discussion as the following corollary.

Corollary 1. (*Limit Laws of LR and W statistics under Condition C*). *Under Condition C, Conditions R holds with the limit Gaussian process stated in equation (2.7). The plug-in estimates of the weighting functions (2.12) and (2.13) are uniformly consistent for the weighting functions (2.10) and (2.11). Therefore, conclusions of Theorem 1 hold with the limit laws for our statistics given by the laws of random variables stated in equations (2.8) and (2.9).*

2.3. Basic Validity and Convergence Rates for the Confidence Regions. We will first give a basic validity result for confidence regions assuming that we have suitable estimates of the quantiles of LR and W statistics and will verify basic validity of our confidence regions. A basic procedure for constructing suitable estimates of these quantiles via bootstrap or simulation will be given below.

Theorem 2. (*Basic Inferential Validity of Confidence Regions*). *Suppose that for $\alpha < 1/2$ we have consistent estimates of quantiles of limit statistics \mathcal{W} and \mathcal{L} , namely,*

$$\widehat{k}(1 - \alpha) = k(1 - \alpha) + o_p(1), \quad (2.14)$$

where $k(1 - \alpha)$ is the $(1 - \alpha)$ -quantile of either \mathcal{W} or \mathcal{L} , respectively. Then as the sample size n grows to infinity, confidence regions R_{LR} and R_W cover Θ_0 with probability approaching $1 - \alpha$:

$$P[\Theta_0 \subseteq R_{LR}] = P[\mathcal{L}_n \leq \widehat{k}(1 - \alpha)] \rightarrow P[\mathcal{L} \leq k(1 - \alpha)] = (1 - \alpha), \quad (2.15)$$

$$P[\Theta_0 \subseteq R_W] = P[\mathcal{W}_n \leq \widehat{k}(1 - \alpha)] \rightarrow P[\mathcal{W} \leq k(1 - \alpha)] = (1 - \alpha). \quad (2.16)$$

We next recall that given a Euclidian metric $d(\cdot, \cdot)$ on \mathbb{R}^K , the (symmetric) Hausdorff distance between two non-empty sets $A, B \subset \mathbb{R}^k$ is defined as

$$d_H(A, B) := \max \left\{ \sup_{a \in A} d(a, B), \sup_{b \in B} d(b, A) \right\}. \quad (2.17)$$

Our next result shows that the confidence regions based on the LR and the Wald statistic are also root- n consistent estimators for the set Θ_0 with respect to the Hausdorff distance. This result also establishes that the Pitman rates for R_W and R_{LR} is also root- n , namely any alternative set Θ_A such that $\sqrt{n}d_H(\Theta_0, \Theta_A) \rightarrow \infty$ will not be covered by R_W and R_{LR} , and hence will be rejected with probability tending to 1.

Theorem 3. (*Confidence Regions are \sqrt{n} -Consistent Estimator of Θ_0*). *Under Condition R, the confidence regions based on the LR and Wald statistics are consistent with respect to the Hausdorff distance at a root- n rate, that is*

$$d_H(R_{LR}, \Theta_0) = O_p(n^{-1/2}) \text{ and } d_H(R_W, \Theta_0) = O_p(n^{-1/2}).$$

As a consequence, if $\sqrt{n}d_H(\Theta_0, \Theta_A) \rightarrow \infty$, then

$$P[\Theta_A \subseteq R_{LR}] \rightarrow 0 \text{ and } P[\Theta_A \subseteq R_W] \rightarrow 0.$$

While the results in Theorems 1-3 were stated in terms of the high-level Condition R, the next corollary shows that these results can be applied to any problem where Condition C holds. These conditions hold in HJ, MF, and other problems listed in Section 1, as we verify below.

Corollary 2. (*Limit Laws of LR and W statistics under Condition C*). *Suppose that Condition C holds and that consistent estimates of quantiles of statistics (2.8) and (2.9) are available. Then the conclusions of Theorem 2 apply.*

We now turn to estimation of quantiles for the LR and W statistics using bootstrap, simulation, and other resampling schemes under general conditions. The basic idea is as follows: First, let us take any procedure that consistently estimates the law of our basic Gaussian process \mathbb{G} or a weighted version of this process appearing in the limit expressions. Next, we can use the estimated law to obtain consistent estimates of the laws of the LR and W statistics, and thus also obtain consistent estimates of their quantiles. It is well known that there are many procedures for accomplishing the first step, including such common schemes as the bootstrap, simulation, and subsampling, which can also be adapted to allow for various forms of cross-section and time series dependence.

In what follows, we will simplify the notation by writing our limit statistics as a special case of the following statistic:

$$\mathcal{S} := \sup_{\theta \in \partial\Theta_0} [V(\theta)]_+^2, \quad V(\theta) := \tau(\theta)\mathbb{G}(\theta). \quad (2.18)$$

Thus, $\mathcal{S} = \mathcal{L}$ for $\tau(\theta) = 1/s(\theta)$ and $\mathcal{S} = \mathcal{W}$ for $\tau(\theta) = 1/[\|\nabla_{\theta}m(\theta)\| \cdot w(\theta)]$. We take τ to be a continuous function bounded away from zero on the parameter space. We also need to introduce the following notations and concepts. Our process V is a random element that takes values in the metric space of continuous functions $C(\Theta)$ equipped with the uniform metric. The underlying measure space is (Ω, \mathcal{F}) and we denote the law of V under the probability measure P by the symbol \mathcal{Q}_V .

In the following we will assume that we have an estimate \mathcal{Q}_{V^*} of the law \mathcal{Q}_V of the Gaussian process V . This estimate \mathcal{Q}_{V^*} is a probability measure which can be generated as follows: Let us fix another measure space (Ω', \mathcal{F}') and a probability measure P^* on this space. Then given a random element V^* on this space taking values in $C(\Theta)$, we denote its law under P^* by \mathcal{Q}_{V^*} . We thus identify the probability measure P^*

with a data-generating process by which we generate draws or realizations of V^* . This identification allows us to cover such methods of producing realizations of V^* as the bootstrap, subsampling, or other simulation approaches.

We require that the estimate \mathcal{Q}_{V^*} be consistent for \mathcal{Q}_V in any metric ρ_K metrizing weak convergence, where we can take the metric to be the Kantarovich-Rubinstein metric. Note that there are many results that verify this basic consistency condition for different processes V and various bootstrap, simulation, and subsampling schemes, as we will discuss in more detail below.

To define the Kantarovich-Rubinstein metric, let $\theta \mapsto v(\theta)$ be an element of a metric space (M, d_M) , and $\text{BL}_1(M)$ be a class of Lipschitz functions $\varphi : M \rightarrow \mathbb{R}$ that satisfy:

$$|\varphi(v) - \varphi(v')| \leq d_M(v, v'), \quad \sup_{v \in M} |\varphi(v)| \leq 1.$$

The Kantarovich-Rubinstein distance between probability laws \mathcal{Q}_V and $\mathcal{Q}_{V'}$ of random elements V and V' taking values in M is defined as:

$$\rho_K(\mathcal{Q}, \mathcal{Q}'; M) := \sup_{\varphi \in \text{BL}_1(M)} |\mathbb{E}_{\mathcal{Q}}\varphi(V) - \mathbb{E}_{\mathcal{Q}'}\varphi(V)|.$$

As stated earlier, we require the estimate \mathcal{Q}_{V^*} to be consistent for \mathcal{Q}_V in the metric ρ_K , that is

$$\rho_K(\mathcal{Q}_{V^*}, \mathcal{Q}_V; C(\Theta)) = o_p(1). \quad (2.19)$$

Let $\mathcal{Q}_{\mathcal{S}}$ denote the probability law of $\mathcal{S} = \mathcal{W}$ or \mathcal{L} , which is in turn induced by the law \mathcal{Q}_V of the Gaussian process V . We need to define the estimate $\mathcal{Q}_{\mathcal{S}^*}$ of this law. First, we define the following plug-in estimate of the boundary set $\partial\Theta_0$,

$$\widehat{\partial\Theta_0} = \{\theta \in \Theta : \widehat{m}(\theta) = 0\}. \quad (2.20)$$

This estimate turns out to be consistent at the usual root- n rate, by an argument similar to that given in Chernozhukov et al. (2007). Next, define $\mathcal{Q}_{\mathcal{S}^*}$ as the law of the random variable

$$\mathcal{S}^* = \sup_{\theta \in \widehat{\partial\Theta_0}} [V^*(\theta)]_+^2. \quad (2.21)$$

In this definition, we hold the hatted quantities fixed, and the only random element is V^* that is drawn according to the law \mathcal{Q}_{V^*} .

We will show that the estimated law $\mathcal{Q}_{\mathcal{S}^*}$ is consistent for $\mathcal{Q}_{\mathcal{S}}$ in the sense that

$$\rho_K(\mathcal{Q}_{\mathcal{S}^*}, \mathcal{Q}_{\mathcal{S}}; \mathbb{R}) = o_p(1). \quad (2.22)$$

Consistency in the Kantorovich-Rubinstein metric in turn implies consistency of the estimates of the distribution function at continuity points, which in turn implies consistency of the estimates of the quantile function.

Equipped with the notations introduced above we can now state our result.

Theorem 4. (*Consistent Estimation of Critical Values* $k(1 - \alpha)$). *Suppose Conditions R.1-R.3 hold, and that we have a consistent estimate of the law of our limit Gaussian processes V . Then the estimates of the laws of the limit statistics $\mathcal{S} = \mathcal{W}$ or \mathcal{L} defined above are consistent. As a consequence, we have that the estimates of the quantiles are consistent in the sense of equation (2.14).*

It is useful to give a similar result under more primitive conditions C.1-C.2. Recall that in this case our estimator satisfies

$$\sqrt{n}(\hat{\gamma} - \gamma) \rightsquigarrow \Omega^{1/2}Z,$$

so that our limit statistics take the form:

$$\mathcal{S} = \sup_{\theta \in \partial\Theta_0} [V(\theta)]_+^2, \quad V(\theta) = t(\theta)'Z,$$

where $t(\theta)$ is a vector valued weight function, in particular, we have

$$\begin{aligned} t(\theta) &= (\nabla_\gamma m(\theta, \gamma)' \Omega^{1/2}) / s(\theta) && \text{for } \mathcal{S} = \mathcal{L}, \\ t(\theta) &= (\nabla_\gamma m(\theta, \gamma)' \Omega^{1/2}) / (\|\nabla_\theta m(\theta, \gamma)\| \cdot w(\theta)) && \text{for } \mathcal{S} = \mathcal{W}. \end{aligned}$$

Here we shall assume that we have a consistent estimate \mathcal{Q}_{Z^*} of the law \mathcal{Q}_Z of Z , in the sense that,

$$\rho_K(\mathcal{Q}_{Z^*}, \mathcal{Q}_Z) = o_p(1). \tag{2.23}$$

For instance we can simulate the distribution using draws $Z^* \sim N(0, I)$ or apply any valid bootstrap method to $\hat{\Omega}^{-1/2} \sqrt{n}(\hat{\gamma} - \gamma)$. Alternatively it is possible to use subsampling (Politis and Romano (1994)). Then the estimate \mathcal{Q}_{V^*} of the law \mathcal{Q}_{V^*} is defined as:

$$V^*(\theta) = \hat{t}(\theta)'Z^*, \tag{2.24}$$

where $\hat{t}(\theta)$ is a vector valued weighting function that is uniformly consistent for the weighting function $t(\theta)$. In particular, we can use

$$\begin{aligned} \hat{t}(\theta) &= (\nabla_\gamma m(\theta, \hat{\gamma})' \hat{\Omega}^{1/2}) / \hat{s}(\theta) && \text{for } \mathcal{S} = \mathcal{L}, \\ \hat{t}(\theta) &= (\nabla_\gamma m(\theta, \hat{\gamma})' \hat{\Omega}^{1/2}) / (\|\nabla_\theta m(\theta, \hat{\gamma})\| \cdot w(\theta)) && \text{for } \mathcal{S} = \mathcal{W}. \end{aligned}$$

In this definition we hold the hatted quantity fixed, and the only random element being resampled or simulated is Z^* , with the law denoted as \mathcal{Q}_{Z^*} . Then, we define the random variable

$$\mathcal{S}^* = \sup_{\theta \in \widehat{\partial\Theta}_0} [V^*(\theta)]_+^2,$$

and use its law $\mathcal{Q}_{\mathcal{S}^*}$ to estimate the law $\mathcal{Q}_{\mathcal{S}}$. In the definition of the law $\mathcal{Q}_{\mathcal{S}^*}$ we fix, i.e. condition on, the estimated quantities $\widehat{t}(\cdot)$ and $\widehat{\partial\Theta}_0$.

We can now state the following corollary, which is proven in the appendix:

Corollary 3. (*Consistent Estimation of Critical Values under Condition C*).

Suppose that conditions C holds, and that we have a consistent estimate of the law of Z , so equation (2.23) holds. Then this provides us with a consistent estimate of the law of our limit Gaussian process \mathbb{G} , and hence all of the conclusions of Theorem 3 hold.

2.4. Verification of Condition C for MF- and HJ-Bounds. We can now apply this result to the HJ and MF problems. As before, let $v := \mathbb{E}[r_{t+1}]$ and $\Sigma := \mathbb{E}[(r_{t+1} - v)(r_{t+1} - v)']$, and the parameters $\theta = (\mu, \sigma)'$, and $\gamma_{HJ} = (S_{vv}, S_{v1}, S_{11})'$ for the HJ bounds, where

$$S_{vv} = v'\Sigma^{-1}v, \quad S_{v1} = 1'_N \Sigma^{-1}v, \quad \text{and} \quad S_{11} = 1'_N \Sigma^{-1}1_N.$$

Also let $\widehat{\gamma}_{HJ}$ denote the sample analog obtained by replacing v and Σ with the mean \widehat{v} and variance $\widehat{\Sigma}$ of the sample r_1, \dots, r_T . For the MF bounds, let $\gamma_{MF} = \text{vec}(D'\Sigma D)$, where

$$B := \text{Cov}_t(Z_{t+1}, r_{t+1}) \quad \text{and} \quad D := [v, B']$$

and $\widehat{\gamma}_{MF}$ be the sample analog, and let μ^* and β^* be the target mean portfolio return and factor loadings.

Recall that the Hansen-Jagannathan mean-variance bound for the stochastic discount factor is defined by the moment condition

$$m_{HJ}(\theta, \gamma_{HJ}) = \sqrt{S_{vv}\mu^2 - 2S_{v1}\mu + S_{11}} - \sigma.$$

whereas the multi-factor efficient (MF) mean-variance set for portfolio returns is characterized by

$$m_{MF}(\theta, \gamma_{MF}) := \sqrt{(\mu^*, \beta^*)(D'\Sigma D)^{-1}(\mu^*, \beta^*)'} - \sigma$$

We now give regularity conditions for the validity of inference and confidence intervals based on the LR and Wald statistics derived from the moment functions $\widehat{m}_{HJ}(\theta)$ and $\widehat{m}_{MF}(\theta)$, respectively.

Proposition 1 (Verification of Condition C for HJ and MF problems). *Suppose that $\widehat{v}, \widehat{\Sigma}$, and \widehat{B} satisfy a CLT. Furthermore, let Θ be a rectangle in $\mathbb{R} \times \mathbb{R}_+$, that the absolute values of the elements of v and eigenvalues of Σ are bounded between finite strictly positive constants, and that $N\sum_{i=1}^N v_i^2 - \left(\sum_{i=1}^N v_i\right)^2$ is bounded away from zero. Then Condition C holds for the HJ and MF bounds.*

For the last condition, note that for the purposes of this paper, we treat the number of assets N as finite, and that the Cauchy-Schwarz inequality implies that the difference $N\sum_{i=1}^N v_i^2 - \left(\sum_{i=1}^N v_i\right)^2 \geq 0$. The difference will be strictly positive only if the mean return v_i is not constant across assets.

2.5. Structured Projection Approach. We also compare our procedure to alternative confidence regions from a structured projection approach that is based on a confidence set for the point-identified parameter γ that characterizes the bound $m(\theta) = m(\theta, \gamma)$ under the condition C. This confidence region is then “projected” to obtain a confidence region for Θ_0 . We call the projection approach “structured” when γ is a minimal sufficient parameter for the bounds on θ , and $\widehat{\gamma}$ is the minimal sufficient statistics. For example the HJ bounds from Example 1 can be characterized by the three-dimensional parameter

$$\gamma = (S_{vv}, S_{v1}, S_{11})',$$

whereas the Markowitz-Fama mean-variance frontiers are described in terms of the first two moments of the spanning tangency portfolios. Note that the structured approach avoids creating some confidence regions for high-dimensional mean and variance parameters v and Σ and then projecting them to obtain a confidence region for Θ_0 . Such an approach would be extremely conservative, and working with the minimal parameter γ instead reduces the conservativeness dramatically. Since inference for γ is standard, the structured projection approach seems much more natural and “economically appealing”. However, despite its dimension-reducing and intuitive appeal, this approach remains very conservative and is much less powerful than the approach based on the optimally weighted LR-type and W-statistics. We illustrate the superior performance of LR-based confidence sets compared to the structured projection approach empirically for the Hansen-Jagannathan set (see Figure 7).

For the construction of projection confidence sets, we assume that we can construct a $(1 - \alpha)$ confidence set R_γ for γ_0 . This confidence region can then be projected onto the parameter space Θ to form the confidence set for Θ_0 :

$$R_{Proj} := \bigcup_{\gamma \in R_\gamma} \{\theta \in \Theta : m(\theta, \gamma) \leq 0\}. \quad (2.25)$$

Under condition C, $\sqrt{n}(\hat{\gamma} - \gamma_0)$ satisfies a CLT, so we can construct an elliptical joint $1 - \alpha$ confidence region for the quantity γ as follows:

$$R_\gamma = \{\gamma \in \Gamma : n(\hat{\gamma} - \gamma)' \hat{\Omega}^{-1}(\hat{\gamma} - \gamma) \leq \hat{c}(1 - \alpha)\}, \quad (2.26)$$

where $\hat{c}(1 - \alpha)$ is either $1 - \alpha$ -quantile of $\chi^2(k)$ variable, or any consistent estimate of such a quantile. Under asymptotic normality of $\hat{\gamma}$, this construction approximates an upper contour set of the density of the estimator, and therefore gives an (approximate) smallest-volume confidence set for the parameter γ .

Proposition 2. (*Basic Validity of Structured Projection Approach*). *Let R_γ be $1 - \alpha$ confidence set for the parameter γ . Then $\mathbb{P}[\Theta_0 \subset R_{proj}] \geq \mathbb{P}[\gamma_0 \in R_\gamma] = 1 - \alpha$. In particular, under Condition C, the region R_γ in (2.26) obeys*

$$\mathbb{P}[\gamma_0 \in \Gamma] = 1 - \alpha + o(1),$$

so that the confidence region (2.25) based on structured projection obeys

$$\mathbb{P}[\Theta_0 \subset R_{proj}] \geq 1 - \alpha + o(1).$$

Projected confidence sets were first proposed by Scheffé (1953) for the problem of joint confidence bounds for all linear combinations of the form $c'\gamma$, where $\gamma \in \mathbb{R}^d$ is a parameter vector and $c \in \mathbb{S}^{d-1}$, the $d - 1$ dimensional unit sphere in \mathbb{R}^d . For the problem of confidence bands for the linear regression function, it has been shown that optimality of Scheffé (1953)'s method depends crucially on equivariance with respect to translations and orthogonal transformations of the original parameter, see sections 9.4 and 9.5 in Lehmann and Romano (2005) for a discussion. Bohrer (1973) showed optimality of projection bounds for a generalization of Scheffé (1953)'s original problem, but they were shown to be suboptimal if coverage was only required for a restricted set of functionals corresponding to vectors c for a proper subset of \mathbb{S}^{d-1} , see Casella and Strawderman (1980) and Naiman (1984).

In our case, there are at least three regards in which equivariance with respect to the reduced-form parameters γ fails: in all our examples, (1) the dimension of θ is strictly

lower than that of γ , (2) we consider inference problems that are one-sided rather than symmetric where the parameter space is only a compact subset of \mathbb{R}^k , and (3) the inequality generating function $m(\theta, \gamma)$ is nonlinear in γ . Even after approximate linearization, our inference problem reduces to inference on $c'\gamma$ with c restricted to a small subset of the sphere \mathbb{S}^{d-1} , which is exactly the case where the structured projection approach becomes suboptimal. As mentioned above, we illustrate the superior performance of LR-based confidence set compared to the structured projection based set empirically for the HJ problem (see Figure 7) – the LR-based set is much smaller and lies strictly inside the projection-based set.

2.6. Invariance and Similarity Properties of Confidence Regions based on LR and Wald Statistics.

We next proceed to state the invariance properties of the proposed inference procedures with respect to parameter transformations. We distinguish between an exact invariance and an asymptotic invariance. A parameter transformation is a one-to-one mapping $\eta : \Theta \rightarrow \Upsilon$, where $\Upsilon = \eta(\Theta)$, and we denote the population and sample moment conditions for the transformed problem by $m_\eta(\cdot) := m(\eta^{-1}(\cdot))$, and $\widehat{m}_\eta(\cdot) := \widehat{m}(\eta^{-1}(\cdot))$, respectively, which are mappings from $\Upsilon \rightarrow \mathbb{R}$. We define H^* as the set of parameter transformations η that are continuously differentiable in $\theta \in \Theta^\delta$ and such that Conditions R.1-R.3 hold for the transformed moments $m_\eta(\cdot)$.

In this section we discuss invariance properties of inference based on the LR and W statistics. Let

$$\mathcal{L}_n(\theta) := [\sqrt{n}\widehat{m}_n(\theta)/\widehat{s}(\theta)]_+^2 \quad \text{and} \quad \mathcal{W}_n(\theta) := \left(\sqrt{nd}(\theta, \widehat{\Theta}_{0,\delta})/\widehat{w}(\theta) \right)^2,$$

and let the critical values $\widehat{k}_L(1 - \alpha)$ and $\widehat{k}_W(1 - \alpha)$ be consistent estimators of the asymptotic $1 - \alpha$ quantiles of $\mathcal{L}_n = \sup_{\theta \in \Theta_0} \mathcal{L}_n(\theta)$ and $\mathcal{W}_n = \sup_{\theta \in \Theta_0} \mathcal{W}_n(\theta)$, respectively. Consider the decision functions for including a value of θ in the confidence regions based on the LR and Wald-type statistics, respectively,

$$\phi_{L_n}(\theta) := \mathbf{1}\{\mathcal{L}_n(\theta) \leq \widehat{k}_L(1 - \alpha)\} \quad \text{and} \quad \phi_{W_n}(\theta) := \mathbf{1}\{\mathcal{W}_n(\theta) \leq \widehat{k}_W(1 - \alpha)\},$$

respectively. Similarly, given the parameter transformation η , we consider the statistics

$$\mathcal{L}_n(\theta; \eta) := [\sqrt{n}\widehat{m}_\eta(\eta(\theta))/\widehat{s}(\theta)]_+^2 \quad \text{and} \quad \mathcal{W}_n(\theta; \eta) := \left(\sqrt{nd}(\eta(\theta), \eta(\widehat{\Theta}_{0,\delta}))/\widehat{w}(\theta; \eta) \right)^2$$

with a weight function $\widehat{w}(\theta; \eta)$ possibly depending on η , and the resulting decision functions

$$\phi_{L_n}(\theta; \eta) := \mathbf{1}\{\mathcal{L}_n(\theta; \eta) \leq \widehat{k}_{L_n}(1 - \alpha)\} \quad \text{and} \quad \phi_{W_n}(\theta; \eta) := \mathbf{1}\{\mathcal{W}_n(\theta; \eta) \leq \widehat{k}_{W_n}(1 - \alpha)\},$$

where the critical values $\widehat{k}_{L_n}(1 - \alpha)$ and $\widehat{k}_{W_n}(1 - \alpha)$ are estimates of the asymptotic $1 - \alpha$ quantiles of $\mathcal{L}_n(\eta) = \sup_{\theta \in \partial\Theta_0} \mathcal{L}_n(\theta; \eta)$ and $\mathcal{W}_n(\eta) = \sup_{\theta \in \partial\Theta_0} \mathcal{W}_n(\theta; \eta)$, respectively.

We say that the decision function $\phi_n(\theta; \eta)$ is

- *invariant* if for any $\theta \in \Theta$, we have $\phi_n(\theta) = \phi_n(\theta; \eta)$ for any parameter transformation $\eta \in H^*$,
- *asymptotically invariant* to first order if for any parameter transformation $\eta \in H^*$, any $\theta_0 \in \Theta$, and any sequence $\theta_n \rightarrow \theta$, we have $P(\phi_n(\theta_n) \neq \phi_n(\theta_n; \eta)) \rightarrow 0$ as $n \rightarrow \infty$.

These invariance properties describe whether the parameter transformations affect the inclusion of any sequence of points θ_n in the confidence sets. Note in particular that (exact) invariance is a property of a given realization of $\phi_n(\theta; \eta)$ and implies asymptotic invariance. Furthermore, it is easy to verify that invariance of a decision function $\phi_n(\theta; \eta)$ implies analogous equivariance properties for the corresponding confidence sets

$$R_{LR} \equiv \{\theta \in \Theta : \phi_{L_n}(\theta; \eta) = 1\} \quad \text{and} \quad R_W \equiv \{\theta \in \Theta : \phi_{W_n}(\theta; \eta) = 1\}.$$

We also would like to mention another property, which characterizes the precision of the confidence sets. For a given value of η , we also say that a test based on $\phi_n(\theta; \eta)$ is *asymptotically similar* on the boundary of Θ_0 if $\lim_n E[\phi_n(\theta; \eta)] = c$ for some constant $c \in (0, 1)$ and any $\theta \in \partial\Theta_0$. Similarity here means that any point on the boundary of the identified set can be expected to be included with asymptotic probability c , which does not vary with the location of the point.

Given these definitions, we can now characterize the invariance and similarity properties of inference procedures based on the weighted LR and Wald statistics:

Proposition 3. (*Invariance Properties of Decision Functions based on $\mathcal{L}_n(\theta; \eta)$ and $\mathcal{W}_n(\theta; \eta)$*). *Suppose Conditions R.1-R.3 hold. Then (i) the decision function $\phi_{L_n}(\theta; \eta)$ based on $\mathcal{L}_n(\theta; \eta)$ is invariant, and asymptotically similar on the boundary of Θ_0 , whereas (ii) the decision function $\phi_{W_n}(\theta; \eta)$ based on $\mathcal{W}_n(\theta; \eta)$ is asymptotically invariant to first order if the weighting function is of the form $\widehat{w}(\theta; \eta) = \frac{b(\theta)}{\|\nabla_{\eta} m_{\eta}(\eta(\theta))\|} + o_p(1)$ for some function $b(\theta) > 0$ that does not depend on η and all $\theta \in \partial\Theta$. Furthermore, (iii)*

if $\widehat{w}(\theta; \eta) = b \frac{\widehat{s}(\theta)}{\|\nabla_{\eta} m_{\eta}(\eta(\theta))\|} + o_p(1)$ for some constant $b > 0$, then $\phi_{W_n}(\theta)$ is asymptotically similar on the boundary of Θ_0 .

Notice in particular the different roles the norm of the gradient of $m(\theta)$ and the standard deviation $s(\theta)$ play for the properties of the weighted Wald statistic: Choosing weights that are inversely proportional to $\|\nabla_{\theta} m(\theta)\|$ in the limit corrects for the dependence of the Hausdorff distance on the parameterization of the problem, and accounting for $s(\theta)$ also gives similarity on the boundary. In particular, only the weights in part (iii) of Proposition 3 yield results for confidence sets based on the Wald statistic that compare to the performance of LR-based inference. Our empirical results below illustrate that the difference is important, since the lack of invariance or precision can lead to overturning the main economic conclusions in the empirical analysis.

3. EMPIRICAL APPLICATIONS

3.1. Hansen-Jagannathan Mean-Variance Sets for the SDF. In order to illustrate the performance of our procedure, we estimate confidence sets for the Hansen-Jagannathan sets of mean-variances of stochastic discount factors. In order to keep results comparable with Hansen and Jagannathan (1991), we construct the sample for the empirical exercise following the data description in Hansen and Jagannathan (1991). The two asset series used are annual treasury bond returns and annual NYSE value-weighted dividend included returns. These nominal returns are converted to real returns by using the implicit price deflator based on personal consumption expenditures used by Hansen and Jagannathan (1991). Asset returns are from CRSP, and the implicit price deflator is available from St. Louis Fed and based on National Income and Product Accounts of United States. We use data for the years 1959-2006.

Figure 1 reports the estimated bound consisting of the mean-standard deviation pairs which satisfy

$$m(\theta, \widehat{\gamma}) = 0,$$

where $\widehat{\gamma}$ is estimated using sample moments.

We can compare the estimated HJ bounds with mean-variance combinations implied by the consumption CAPM model. In the model the economy is equivalent to a representative agent with constant elasticity of intertemporal substitution preferences

$$u(C_t) = \frac{C_t^{1-\varrho}}{1-\varrho},$$

where C_t is the aggregate consumption. Then the stochastic discount factor implied by consumption growth is given by $M_t(\varrho) = \beta \left(\frac{C_t}{C_{t+1}} \right)^\varrho$, so that we can estimate the first two moments of $M_t(\varrho)$. Specifically, let $\mu_C(\varrho) := \mathbb{E}[M_t(\varrho)]$ and $\sigma_C(\varrho) := \sqrt{\text{Var}(M_t(\varrho))}$ be the mean and standard deviation, respectively, of the stochastic discount factor given an iso-elastic utility function with an elasticity of intertemporal substitution equal to $1/\varrho$. We can characterize the feasible set Θ_e of mean-variance pairs by the moment restriction

$$0 = m((\mu, \sigma)) := \mu_C(\sigma_C^{-1}(\sigma)) - \mu.$$

where $\sigma_C^{-1}(\cdot)$ denotes the inverse function of $\sigma_C(\varrho)$.³ Given an i.i.d. sample of observations for the growth rate of consumption $\frac{C_{t+1}}{C_t}$, we define the empirical analog $\widehat{m}((\mu, \sigma))$ analogously, where the expectations in the definition of $\mu_C(\cdot)$ and $\sigma_C(\cdot)$ are replaced with averages. The mean-variance pairs reported in Figure 1 were obtained using data on per capita expenditures for non-durable consumption and services in the U.S. from 1959-2006, assuming a discount factor $\beta = 0.95$.

It is well known that it is difficult to reconcile asset prices and aggregate consumption empirically in a representative framework. For our data, the values of ϱ corresponding to mean-variance pairs for the SDF that fall inside the estimated HJ set range from about 170 to 192, suggesting a very low elasticity of intertemporal substitution and implying a high variance for the SDF. These values for ϱ are unrealistically large, but are in line with other findings in the empirical literature on the consumption based CAPM.⁴

In order to represent the sampling uncertainty in estimating γ , we plot 100 bootstrap draws of the HJ frontier in Figure 2, where observations were drawn with replacement from the bivariate time series of stock and bond returns. In order to represent the sampling uncertainty in estimating the mean-standard deviation pairs $(\mu_C(\varrho), \sigma_C(\varrho))$ of the consumption-based SDF, we also plot 100 bootstrap draws of $\varrho \mapsto (\widehat{\mu}_C(\varrho), \widehat{\sigma}_C(\varrho))$. We see that the sampling uncertainty is quite considerable for both HJ frontier and for mean-standard deviation pairs implied by the consumption-based SDF. In fact, an intriguing feature of this graph is that near the apex of the HJ frontier, the pair $(\widehat{\mu}_C(\varrho), \widehat{\sigma}_C(\varrho))$ and its bootstrap draws are close to the HJ frontier at low values of ϱ . The low values of ϱ are considered to be "reasonable," since they correspond to a relatively high intertemporal elasticity of substitution and appear to be well micro-founded. The inability of

³It can be verified that $\sigma_C(\varrho)$ is strictly increasing in ϱ , so that this inverse is well-defined.

⁴E.g. in Lettau and Ludvigson (2009), the values for the elasticity of intertemporal substitution minimizing the mean squared error in the Euler equations characterizing the household's investment problem are comparable in magnitude.

“reasonable” values of ϱ to reconcile aggregate consumption data with asset prices has been a major theme of the empirical literature on the consumption-based CAPM model starting with Hansen and Singleton (1982). However from Figures 1 and 2 it is not obvious that an empirical test will reject the canonical/baseline model underlying that literature, and we report confidence sets based on the various approaches discussed in earlier sections based on which we can make inferential statements about the benchmark model.

Figure 3 shows the 95% confidence region based on the LR statistic. By construction, the LR confidence region covers most of the bootstrap draws below the HJ bounds. However, it should also be noted that the confidence bound based on the LR statistic is fairly tight relative to the bootstrapped frontiers, and does not include any unnecessary areas of the parameter space. Noting that the set Θ_e for the consumption-based SDF is defined by a moment equality, we can also form a confidence band R_e for Θ_e based on the statistic $\mathcal{L}_n(\theta) := \left[\frac{\widehat{m}_{SDF}(\theta)}{\widehat{s}(\theta)} \right]^2$, noting that the asymptotic arguments in the derivations for the (one-sided) LR statistic can be easily extended to the two-sided case if we replace squared positive parts $[\cdot]_+^2$ with the usual square, $(\cdot)^2$. The lower and upper bounds in the following figures were constructed using separate estimates of the local standard deviation $s((\mu, \sigma))$ based on the negative and positive deviations of $\widehat{m}((\mu, \sigma))$, respectively to improve the approximation. Critical values were obtained using the nonparametric bootstrap.

Most importantly, the LR-based confidence region for the HJ set does not overlap with the confidence set for the consumption-based SDF for “small” values of ϱ (in fact, for any $\varrho \in [0, 120]$). The absence of overlap for the 95% regions implies the rejection of any $\varrho \in [0, 120]$ at 10% significance level. This is clear evidence against the benchmark formulation of the consumption-based CAPM, and therefore an important empirical conclusion. In what follows below we will show that the same empirical conclusion cannot be reached for this example using less precise or non-invariant methods. Specifically, we will show that if we use confidence regions based on either LR-statistic without precision weighting, or Wald statistics without invariance/precision weighting, or regions based on structural projection, we will not be able to reach the same empirical conclusion. So invariance and precision considerations in construction of the confidence regions turn out to be quite important for reaching sharp economic conclusions.

Figure 4 plots the 95% confidence region based on an unweighted LR statistic. Comparing Figure 3 and Figure 4 it can be seen that precision weighting plays a very important role in delivering good confidence sets. Without precision weighting, the unweighted LR statistic delivers a confidence region that includes implausible regions in the parameter space where the standard deviation of the discount factor is zero. Moreover, the confidence region becomes too imprecise to reject the canonical model of the stochastic discount factor.

Figure 6 plots the confidence region based on the Wald statistic with no invariance/precision weighting, which is equivalent to a confidence region based on the directed Hausdorff distance. Similar to Figure 4 the confidence set covers a large area of the parameter space which is excluded from any bootstrap realization of the HJ set. The shape of the confidence region based on the Wald statistic in Figure 6 seem counter-intuitive because at first sight, as the confidence bounds do not appear to be a uniform enlargement of the estimated frontier $\partial\hat{\Theta}$. However, this visual impression is only due to the fact that the plot shows units of μ and σ at different scales. The observation that the weighting and scaling of the different components of θ seem “unnatural” in this particular graph emphasizes the potential problems associated with the non-invariance of inference based on the unweighted Wald statistic.

Figure 5 plots the confidence region based on the weighted Wald statistic, where weights induce first order invariance and similarity via precision weighting. This weighting fixes the problem and generates a statistic that is (first-order) invariant to parameter transformations. As a result, the confidence set looks very similar to weighted LR based confidence set in Figure 3 in that it covers most of the bootstrap draws below the HJ bounds and its shape reflects local sampling uncertainty in an adequate manner. This practical evidence therefore emphasizes the importance of introducing invariance and precision inducing weights in the Wald-based approach, which we had argued for theoretically in the previous sections.

Finally, in Figure 7 we compare our results to confidence regions from the structured projection approach that is based on a confidence set for the point-identified parameters. As described in section 2.5, we construct an elliptical joint $1 - \alpha$ confidence region for the quantity γ defined in equation (1.2) based on the quadratic form for the estimator $\hat{\gamma}$, $\hat{T}(\gamma) := (\hat{\gamma} - \gamma)' \text{var}(\hat{\gamma})^{-1} (\hat{\gamma} - \gamma)$. For the diameter of this confidence ellipsoid we used both a bootstrap and a chi-square approximation to the distribution of $\hat{T}(\gamma)$, which both yield qualitatively similar results.

The $1 - \alpha$ confidence set for $\theta = (\mu, \sigma)$ is obtained by projecting the $1 - \alpha$ confidence region for γ onto Θ using the condition $m(\theta, \gamma) \leq 0$. We report the resulting confidence set from the structured projection approach in Figure 7 together with the LR-based confidence set proposed in this paper. The structured projection confidence set performs quite poorly relative to the LR-based confidence set: in particular the latter is much smaller and lies strictly inside the former. In fact, the precision of the confidence set based on structured projection is poor enough to overturn the major empirical conclusion that the consumption-based CAPM cannot be reconciled with small values of ρ .

This should be expected since the projection confidence bounds are based on a confidence set for the point-identified parameter that does not account for the specific shape of the bounds as a function of $\hat{\gamma}$. More specifically, the elliptical joint confidence set for γ (which minimizes volume under joint normality of $\hat{\gamma}$) guards us against deviations from the true value in any direction in \mathbb{R}^3 , but most of these deviations are irrelevant for the bounds for θ , since these are only one-sided and the parameter space for θ is only two-dimensional. The fact that the standard confidence set for γ treats all directions in the parameter space Γ symmetrically may be far from ideal for inference on the (μ, σ) -frontier, since the bound on the standard deviation is a nonlinear function whose derivative with respect to γ varies widely across different values of $(\mu, \sigma(\mu))$. Note that for confidence sets for a point-identified parameter, by the delta method the effect of nonlinearities is asymptotically negligible to first order. However when the object of interest is a set with a nontrivial diameter, the resulting effect is of first order even for large samples.

3.2. Bounds on the Elasticity of Labor Supply. In his meta-analysis, Chetty (2012) reports point-wise confidence bounds for the structural Hicksian elasticity ε of labor supply at the intensive margin for given values of the optimization friction δ . The reported bounds result from the intersection of bounds of the form (1.4) from estimates $\hat{\varepsilon}_j$ obtained from J empirical studies exploiting different natural experiments varying the effective income tax $\Delta_j \log p$, $j = 1, \dots, J$.

We apply the bootstrap procedure proposed in this paper to obtain joint confidence sets for (δ, ε) based on the LR and Wald statistics. More specifically, we consider the moments obtained from individual empirical elasticities $\hat{\varepsilon}_j$

$$\hat{g}_{OF,j}((\varepsilon, \delta)') = g_{OF,j}((\varepsilon, \delta)', \hat{\varepsilon}_j) := \frac{(\varepsilon - \hat{\varepsilon}_j)^2 (\Delta_j \log p)^2}{8\varepsilon} - \delta.$$

These “raw” moments are then aggregated by a smooth function

$$\widehat{m}_{OF}^*(\varepsilon, \delta; \lambda) = \sum_{j=1}^J \frac{\exp(\lambda \widehat{g}_{OF,j}(\varepsilon, \delta))}{\sum_{l=1}^J \exp(\lambda \widehat{g}_{OF,l}(\varepsilon, \delta))} \widehat{g}_{OF,j}(\varepsilon, \delta),$$

where $\lambda \in \mathbb{R}$ is a fixed, positive scalar. Note that as discussed in Section 2, this transformation approximates the maximum of $\widehat{g}_{OF,1}(\varepsilon, \delta), \dots, \widehat{g}_{OF,J}(\varepsilon, \delta)$ as $\lambda \rightarrow \infty$, but satisfies the smoothness conditions for our procedure for any finite value of $\lambda > 0$.

We use a parametric bootstrap to obtain the critical value $k(1 - \alpha)$, where we approximate the sampling distribution of the estimators for the respective elasticities by a joint normal distribution centered around the estimates reported in Panel A of table 1 with standard deviations equal to the respective standard errors and zero covariances. This approach can be justified by an assumption that the studies were based on mutually independent random samples from possibly different populations.

Figure 8 shows that the estimated bounds coincide with the set reported in Figure 8 of Chetty (2012) except for the use of the smoothed maximum function instead of the intersection of (ε, δ) -sets which leads to a slightly wider set. The 95% confidence set based on the LR statistic⁵ reported in Figure 9 is fairly narrow around the estimated bound, and does not appear to differ very much from the the collections of confidence intervals in Chetty (2012). Chetty presents confidence intervals that are pointwise with respect to δ , that is for each fixed value of friction δ , the interval covers structural elasticity ε with a prescribed probability. In contrast, our set estimator covers all plausible values of (δ, ε) with a prescribed probability. Thus it simultaneously performs inference on both structural elasticity ε and the friction amount δ . The LR confidence region is a valid joint confidence set for (ε, δ) and not only point-wise in δ . Furthermore, it is not conservative in that we assume a joint sampling distribution for the elasticity estimates instead of constructing Fréchet-Hoeffding bounds. The LR confidence region excludes all points with $\delta \leq 0.3\%$, so that an optimization friction of at least that size would be needed to reconcile the different elasticities found in the studies considered in this meta-analysis.

⁵ Note that for the LR statistic we used the standard error of the negative part of $\widehat{m}^*(\theta)$ as a weighting function which improves the local approximation due to the asymmetry of the distribution for small values of ϱ . Note that for the $N(0, s(\theta)^2)$ distribution, the standard deviation of the negative part is proportional to $s(\theta)$, so that this weighting scheme is asymptotically equivalent to weighting by the (inverse of the) local standard deviation of $\widehat{m}^*(\theta)$.

Finally, we also report a 95% confidence region based on the Wald statistic *without* optimal re-weighting.⁶ As in the case of HJ bounds, the shape of the resulting confidence set does not reflect the sampling variation in the estimated bounds, and the critical value for the Wald statistic is determined by perturbations of the frontier at very low values for ε . More importantly, in contrast to the LR-based region, the confidence set based on the Wald statistic includes points with $\delta = 0$, failing to reject that the empirical elasticities can be reconciled in a model with no optimization frictions and changing one of the main conclusions of the analysis. Using weighted W statistics instead fixes this problem and gives a confidence set that is very similar to the LR-based confidence set; we do not report this confidence set for brevity.

4. CONCLUSION

In this paper we provide new methods for inference on parameter sets and frontiers that can be characterized by a smooth nonlinear inequality. The proposed procedures are straightforward to implement computationally and have favorable statistical properties. By analyzing the geometric and statistical properties of different statistics, we illustrate the importance of equivariance and similarity considerations for achieving tight confidence regions. In particular, while local weighting is irrelevant for the statistical properties of the estimated frontier, it matters greatly for the size and shape of confidence sets. We also consider smoothed intersection bounds from multiple inequality restrictions, where we give an exact upper bound for the approximation error that depends only on the smoothing parameter.

We illustrate the practical usefulness of these procedures in financial econometrics with various classical examples from mean-variance analysis, including inference on Hansen-Jagannathan mean-variance sets of admissible stochastic discount factors, Markowitz-Fama mean-variance sets of admissible portfolios, and factor-based asset pricing. As a second application, we consider Chetty (2012)'s joint bounds for the elasticity of labor supply and an optimization friction. This example suggests a broader range of uses for set inference in the context of possibly misspecified or incomplete economic models.

In both examples, using invariant or precision-weighted statistics is important for maintaining major empirical conclusions that have been reached informally in prior

⁶In order to adjust for the differences in order of magnitude we constructed the Hausdorff-distance based on the norm $\|(\varepsilon, \delta)\| = \sqrt{\varepsilon^2 + 100 \cdot \delta^2}$. Note that in the graph the confidence region looks poorly centered around the estimated bound, but this optical impression is in fact due to the different scaling of the two axis and the difference in the slope of the frontier above and below its apex.

empirical work, e.g. the inability of large values of the elasticity of intertemporal substitution to generate plausible distributions of stochastic discount factors, or the need for nontrivial optimization frictions to reconcile estimated demand elasticities from different settings. Therefore, the empirical examples illustrate our formal points about the advantages of inference based on a precision weighted metric that is invariant to parameter transformations.

APPENDIX A. PROOFS

Proof of Lemma 1. W.l.o.g., let $\max_j g_j = g_1$ and rewrite

$$\begin{aligned} \lambda \left(\max_j g_j - m(g_1, \dots, g_J; \lambda) \right) &= \frac{\sum_{j=1}^J \exp\{\lambda(g_j - g_1)\} \lambda(g_1 - g_j)}{\sum_{j=1}^J \exp\{\lambda(g_j - g_1)\}} \\ &= \frac{\sum_{j=1}^J \exp\{-h_j\} h_j}{\sum_{j=1}^J \exp\{-h_j\}}, \end{aligned}$$

where $h_j := \lambda(g_1 - g_j) \geq 0$. Clearly, this expression is nonnegative, and since $h_1 = 0$, the denominator is bounded from below by 1. Next note that the function

$$F(h_2, \dots, h_J) := \frac{\sum_{j=2}^J \exp\{-h_j\} h_j}{1 + \sum_{j=2}^J \exp\{-h_j\}}$$

is strictly quasi-concave on \mathbb{R}_+^{J-1} , so that the usual first-order conditions for a local extremum are sufficient for a global maximum. We can now verify that the first-order conditions for maximization of $F(h_2, \dots, h_J)$ have the symmetric solution $h_2 = \dots = h_J = h^* := 1 + W^*$ where $W^* := W\left(\frac{J-1}{e}\right)$. Note that by definition of the product logarithm, $W^* = \frac{J-1}{e} \exp\{-W^*\} = (J-1) \exp\{-1 - W^*\}$, so that

$$\begin{aligned} \max_{h_2, \dots, h_J \geq 0} F(h_2, \dots, h_J) &\equiv \frac{\sum_{j=2}^J \exp\{-h^*\} h^*}{1 + \sum_{j=2}^J \exp\{-h^*\}} = \frac{(J-1) \exp\{-1 - W^*\} (1 + W^*)}{1 + (J-1) \exp\{-1 - W^*\}} \\ &= \frac{W^* + (J-1) \exp\{-1 - W^*\} W^*}{1 + (J-1) \exp\{-1 - W^*\}} = W^* = W\left(\frac{J-1}{e}\right) \end{aligned}$$

Since $\lambda(\max_j g_j - m(g_1, \dots, g_J; \lambda)) = F(\lambda(g_1 - g_2), \dots, \lambda(g_1 - g_J))$, we therefore have that

$$\sup_{g_1, \dots, g_J} \lambda \left| \max_j g_j - m(g_1, \dots, g_J; \lambda) \right| = \sup_{h_2, \dots, h_J \geq 0} |F(h_2, \dots, h_J)| = W\left(\frac{J-1}{e}\right)$$

which establishes the conclusion.

Next, we will prove four lemmas which will be used to justify the local approximation for the Wald statistic. We consider a (stochastic or deterministic) sequence of moment functions $m_n(\theta) := m(\theta) - q_n$, where $q_n = o_p(1)$, and $m(\theta)$ satisfies Conditions R.1-R.2 from the main text, and the corresponding sequence of parameter sets $\Theta_n := \{\theta \in \Theta : m_n(\theta) \leq 0\}$.

Lemma 2. *Suppose the parameter space Θ is compact. Suppose that the gradient $\nabla_{\theta} m_n(\theta)$ is bounded away from zero uniformly in θ and $n = 1, 2, \dots$, and Lipschitz-continuous in θ with Lipschitz constant $L < \infty$. Also let θ_n be any sequence such that θ_n approaches the boundary Θ_n , i.e. $d(\theta_n, \partial\Theta_n) \rightarrow 0$. Then there exists $\delta > 0$ such that the projection of θ_n on Θ_n is unique for all such sequences whenever $d(\theta_n, \Theta_n) < \delta$.*

PROOF: Suppose the statement wasn't true. Then for some sequence θ_n , we could construct a subsequence $\theta_{k(n)}$ such that there are (at least) two distinct projections of $\theta_{k(n)}$ onto $\partial\Theta_n$ for each n . By compactness of Θ , $\theta_{k(n)}$ has a convergent sub-subsequence $\theta_{b(n)}$ with $\lim_n \theta_{b(n)} = \theta_0$, say. Since by construction every member of $\theta_{b(n)}$ has two distinct projections onto $\partial\Theta_n$, we can inscribe a ball of radius $r_n := d(\theta_{b(n)}, \partial\Theta_n)$ centered at $\theta_{b(n)}$ into Θ/Θ_n such that this ball has at least two distinct points $(\theta_{1,b(n)}^*, \theta_{2,b(n)}^*)$ in common with $\partial\Theta_n$.

By properties of the projection, the radii of these balls corresponding to the projection points, $N(\theta_{j,b(n)}^*) := r_n^{-1}(\theta_{b(n)} - \theta_{j,b(n)}^*)$ for $j = 1, 2$, are also normal vectors to the surface $\partial\Theta_n$ at $\theta_{1,b(n)}^*$ and $\theta_{2,b(n)}^*$, respectively. Note that, since the gradient $\nabla_{\theta} m_n(\theta)$ is bounded away from zero, we can w.l.o.g. normalize the length of the normal vectors of the surface $\partial\Theta_n$ to 1.

Note that the two points $\theta_{1,b(n)}^*, \theta_{2,b(n)}^*$ are equidistant to $\theta_{b(n)}$, and therefore lie on a one dimensional sphere $S^1(r_n, \theta_{b(n)}) \subset \Theta$ with center $\theta_{b(n)}$ and radius r_n . The curve $\eta(s)$ corresponding to a sphere of radius r , where s is the arc length, has constant curvature $\kappa(s) \equiv \frac{1}{r}$, so that the normal vectors (with length normalized to one) $N(\theta_{j,b(n)}^*) := r_n^{-1}(\theta_{b(n)} - \theta_{j,b(n)}^*)$ for $j = 1, 2$ satisfy

$$\|N(\theta_{2,b(n)}^*) - N(\theta_{1,b(n)}^*)\| \gtrsim r_n^{-1} \|\theta_{2,b(n)}^* - \theta_{1,b(n)}^*\|.$$

Since $r_n \rightarrow 0$, there is no upper bound on r_n^{-1} , so that $N(\theta)$ is not Lipschitz continuous.

However, the normal vector of $\partial\Theta_0$ at θ standardized to length 1 is given by

$$N(\theta) := \|\nabla_{\theta} m(\theta)\|^{-1} \nabla_{\theta} m(\theta).$$

Condition R.2 implies that $N(\theta)$ is Lipschitz continuous in θ with a constant $1/\delta < \infty$, which leads to a contradiction. Finally, note that since that Lipschitz constant L was assumed to be uniform in γ , there exists $\delta := \delta(L) > 0$ such that projection is unique whenever $r_n := d(\theta_n, \Theta_n) < \delta$ along all sequences θ_n and $m_n(\cdot)$ \square

Lemma 3. *Suppose R.1 and R.2 hold, and let θ_n be a sequence such that $d(\theta_n, \Theta_0) \rightarrow 0$. Define $\bar{\theta}_n \equiv \arg \min_{\widehat{m}(\theta) \leq 0} \|\theta_n - \theta\|^2$. Then $\bar{\theta}_n$ is uniquely defined with probability approaching 1, and satisfies $d(\bar{\theta}_n, \Theta_0) \xrightarrow{p} 0$.*

Proof of Lemma 3. For a given value of $\delta > 0$, let $\eta := \min\{m(\theta) : d(\theta, \Theta_0) \geq \delta\}$. By continuity of $m(\theta)$ and compactness of Θ , $\eta > 0$. By Condition R.1 $\sup_{\theta \in \Theta_0} \widehat{m}(\theta) < \eta$ w.p.a.1, and therefore $d(\bar{\theta}_n, \Theta_0) < \delta$ with probability approaching 1. Uniqueness follows from Lemma 2, noting that the conclusion holds regardless of whether the sequence $m_n(\theta)$ is deterministic or stochastic \square

Lemma 4. *Suppose Conditions R.1 and R.2 hold, and consider a sequence of parameter values $\theta_n \in \Theta$ such that $\theta_n \rightarrow \partial\Theta_0$. For a sequence $q_n = o_p(1)$ and $\widehat{\Theta}_n := \{\theta \in \Theta : \widehat{m}(\theta) \leq q_n\}$, we have*

$$\sqrt{nd}(\theta_n, \widehat{\Theta}_{0,\delta}) = [\|\nabla_{\theta} m(\theta_n)\|^{-1}(\mathbb{G}_n(\theta_n) + \sqrt{n}(m(\theta_n) - q_n)) + o_p(1)]_+$$

In particular, for $q_n \equiv 0$, we have

$$\sqrt{nd}(\theta_n, \widehat{\Theta}_{0,\delta}) = [\|\nabla_{\theta} m(\theta_n)\|^{-1}(\mathbb{G}_n(\theta_n) + \sqrt{nm}(\theta_n)) + o_p(1)]_+$$

Proof of Lemma 4. Without loss of generality, we only consider the case $q_n \equiv 0$. Note that for the general case, we can apply the same argument to the functions $\widehat{m}(\theta) - q_n$ and $m(\theta) - q_n$ which inherit the properties R.1-R.2.

The projection of θ_n onto the set $\widehat{\Theta}_{0,\delta} := \{\theta \in \Theta^\delta : \widehat{m}(\theta) \leq 0\}$ is given by

$$\tilde{\theta}_n := \arg \min_{\theta' \in \Theta^\delta : \widehat{m}(\theta') \leq 0} \|\theta_n - \theta'\|^2.$$

The Lagrangian for this constrained minimization problem is $\|\theta_n - \theta'\|^2 + 2\lambda\widehat{m}(\theta')$, where the Kuhn-Tucker multiplier $\lambda \geq 0$. Since $\delta > 0$ is fixed and $d(\theta_n, \Theta_0) \rightarrow 0$ it follows from Lemma 3 that $\bar{\theta}_n$ is contained in the interior of Θ^δ with probability approaching one. Hence for n large enough, the quantity $\bar{\theta}_n$ can be taken to be an interior solution of the saddle-point problem

$$\begin{aligned} (\bar{\theta}_n - \theta_n) + \nabla_{\theta} \widehat{m}(\bar{\theta}_n) \lambda &= 0, \\ \widehat{m}(\bar{\theta}_n) &= 0. \end{aligned}$$

We can now use a mean-value expansion to obtain

$$\begin{aligned} (\bar{\theta}_n - \theta_n) + \nabla_{\theta} \widehat{m}(\bar{\theta}_n) \lambda &= 0, \\ \widehat{m}(\theta_n) + \nabla_{\theta} \widehat{m}(\bar{\theta}_n)'(\bar{\theta}_n - \theta_n) &= 0, \end{aligned}$$

for an intermediate value $\check{\theta}_n$. By Lemma 3, $\nabla_{\theta}\widehat{m}(\check{\theta}_n) = \nabla_{\theta}m(\theta_n) + o_p(1)$ and $\nabla_{\theta}m(\bar{\theta}_n) = \nabla_{\theta}m(\theta_n) + o_p(1)$. Hence, solving for $(\bar{\theta}_n - \theta_n)$ and applying the extended continuous mapping theorem, we obtain

$$\bar{\theta}_n - \theta_n = -[\nabla_{\theta}m(\theta_n)(\nabla_{\theta}m(\theta_n)'\nabla_{\theta}m(\theta_n))^{-1} + o_p(1)]\widehat{m}(\theta_n)$$

Now, since $\sqrt{n}(\widehat{m} - m) = \mathbb{G}_n = O_p(1)$ in $\ell^\infty(\Theta^\delta)$, we obtain

$$\sqrt{n}(\bar{\theta}_n - \theta_n) = -\nabla_{\theta}m(\theta_n)(\nabla_{\theta}m(\theta_n)'\nabla_{\theta}m(\theta_n))^{-1}(\mathbb{G}_n(\theta_n) + \sqrt{nm}(\theta_n)) + o_p(1)$$

Thus we have that

$$\begin{aligned} \sqrt{nd}(\theta_n, \widehat{\Theta}_n^\delta) &= \sqrt{n}\|\bar{\theta}_n - \theta_n\|1\{\widehat{m}(\theta_n) > 0\} \\ &= |(\nabla_{\theta}m(\theta_n)'\nabla_{\theta}m(\theta_n))^{-1/2}(\mathbb{G}_n(\theta_n) + \sqrt{nm}(\theta_n)) + o_p(1)| \end{aligned} \quad (\text{A.1})$$

$$\times 1\{\mathbb{G}_n(\theta_n) + \sqrt{nm}(\theta_n) > 0 + o_p(1)\} \quad (\text{A.2})$$

$$\begin{aligned} &= [|\nabla_{\theta}m(\theta_n)|^{-1}(\mathbb{G}_n(\theta_n) + \sqrt{nm}(\theta_n)) + o_p(1)]_+ \\ &\quad (\text{A.3}) \end{aligned}$$

which concludes the proof \square

A.1. Proof of Theorem 1. PART 1. (*Limit law of \mathcal{L}_n .*) Let $\mathbb{G}_n = \sqrt{n}(\widehat{m} - m)$. We first show that

$$\begin{aligned} \mathcal{L}_n &= \sup_{\theta \in \Theta_0} [\sqrt{n}\widehat{m}(\theta)/\widehat{s}(\theta)]_+^2 = \sup_{\theta \in \Theta_0} [(\mathbb{G}_n(\theta) + \sqrt{nm}(\theta))/\widehat{s}(\theta)]_+^2 \\ &=_{(1)} \sup_{\theta \in \partial\Theta_0} [(\mathbb{G}_n(\theta) + \sqrt{nm}(\theta))/\widehat{s}(\theta) + o_p(1)]_+^2 \\ &=_{(2)} \sup_{\theta \in \partial\Theta_0} [(\mathbb{G}_n(\theta) + \sqrt{nm}(\theta))/s(\theta) + o_p(1)]_+^2 \\ &\rightsquigarrow_{(3)} \sup_{\theta \in \partial\Theta_0} [(\mathbb{G}(\theta) + \sqrt{nm}(\theta))/s(\theta)]_+^2, \end{aligned}$$

where the steps follows from the following arguments.

To show step (1), note that

$$\sup_{\theta \in \Theta_0} [(\mathbb{G}_n(\theta) + \sqrt{nm}(\theta))/\widehat{s}(\theta)]_+^2 \geq \sup_{\theta \in \partial\Theta_0} [(\mathbb{G}_n(\theta) + \sqrt{nm}(\theta))/\widehat{s}(\theta)]_+^2.$$

So we need to show that with probability approaching 1,

$$\sup_{\theta \in \Theta_0} [(\mathbb{G}_n(\theta) + \sqrt{nm}(\theta))/\widehat{s}(\theta)]_+^2 \leq \sup_{\theta \in \partial\Theta_0} [(\mathbb{G}_n(\theta) + \sqrt{nm}(\theta))/\widehat{s}(\theta) + o_p(1)]_+^2.$$

To show this take a sequence $\theta_n \in \Theta_0$ such that

$$\sup_{\theta \in \Theta_0} [(\mathbb{G}_n(\theta) + \sqrt{n}m(\theta))/\widehat{s}(\theta)]_+^2 = [(\mathbb{G}_n(\theta_n) + \sqrt{n}m(\theta_n))/\widehat{s}(\theta_n)]_+^2,$$

where the supremum is attained by compactness of Θ_0 and continuity hypotheses on \widehat{m} , m , and \widehat{s} , and uniform positivity of \widehat{s} . Since $\sup_{\theta \in \Theta} |\mathbb{G}_n| = O_p(1)$, θ_n can be selected to obey

$$\sqrt{n}m(\theta_n)/\widehat{s}(\theta_n) = O_p(1).$$

We now show that θ_n converges to $\partial\Theta_0$: fix $\kappa > 0$, and let

$$\zeta(\kappa) := \sup \{m(\theta)/s(\theta) : \theta \in \Theta_0, d(\theta, \partial\Theta_0) \geq \kappa\}.$$

Next, note that Condition R.2 implies that $m(\theta) < 0$ for each $\theta \in \text{int } \Theta_0$: suppose this wasn't true, but that $m(\theta_0) = 0$ for some $\theta_0 \in \text{int } \Theta_0$. Then, since for $\theta \in \Theta_0$ $m(\theta) \leq 0$, θ_0 represents a local maximum of $m(\cdot)$ in the interior of Θ_0 . Since $m(\theta)$ is differentiable, this implies that $\nabla_{\theta}m(\theta_0) = 0$. However, Condition R.2 implies that $\|\nabla_{\theta}m(\theta)\| > 0$, a contradiction. Therefore $m(\theta) < 0$ for each $\theta \in \text{int } \Theta_0$, so that by continuity and compactness, $\zeta(\kappa) < 0$.

By conditions R.1 and R.3, we have $\sup_{\theta \in \Theta_0} |\widehat{m}(\theta)/\widehat{s}(\theta) - m(\theta)/s(\theta)| < -\zeta$ with probability approaching one for any strictly negative value of ζ . Hence the supremum of $\sqrt{n}[\widehat{m}(\theta)/\widehat{s}(\theta)]_+^2$ for values $\theta \in \Theta_0$ such that $d(\theta, \partial\Theta_0) \geq \kappa$ is equal to zero with probability approaching one, so that the supremum over all of Θ_0 is attained at a value of θ such that $d(\theta, \partial\Theta_0) < \kappa$. Since κ can be chosen arbitrarily close to zero, we can construct a sequence of positive random variables $\kappa_n = o_p(1)$ such that

$$d(\theta_n, \partial\Theta_0) \leq \kappa_n,$$

with probability approaching 1. We conclude that

$$\begin{aligned} & \sup_{\theta \in \Theta_0} [(\mathbb{G}_n(\theta) + \sqrt{n}m(\theta))/\widehat{s}(\theta)]_+^2 \\ & \leq \sup_{\theta \in \partial\Theta_0, \theta + \eta \in \Theta_0, \|\eta\| \leq \kappa_n} [(\mathbb{G}_n(\theta + \eta) + \sqrt{n}m(\theta + \eta))/\widehat{s}(\theta + \eta)]_+^2. \end{aligned}$$

Using stochastic equicontinuity of \mathbb{G}_n implied by R.1, the last quantity is equal to

$$\sup_{\theta \in \partial\Theta_0, \theta + \eta \in \Theta_0, \|\eta\| \leq \kappa_n} [(\mathbb{G}_n(\theta) + \sqrt{n}m(\theta + \eta))/\widehat{s}(\theta + \eta) + o_p(1)]_+^2,$$

for some $o_p(1)$ term. Because $\sqrt{n}m(\theta + \eta) \leq 0$ for $\theta + \eta \in \Theta_0$ and $m(\theta) = 0$ for $\theta \in \partial\Theta_0$, we conclude that the last quantity is equal to

$$\sup_{\theta \in \partial\Theta_0} [(\mathbb{G}_n(\theta)/\widehat{s}(\theta + \eta) + o_p(1))_+]^2,$$

for the same $o_p(1)$ term. This verifies equality (1).

Equality (2) follows from using R.3 and that the fact that $\sup_{\theta \in \Theta} |\mathbb{G}_n(\theta)| = O_p(1)$ implied by R.1. Equality (3) follows from the application of R.1 and the Continuous Mapping Theorem.

PART 2. (*Limit Law of \mathcal{W}_n*). Recall that we define the set estimator $\widehat{\Theta}_{0,\delta}$ as

$$\widehat{\Theta}_{0,\delta} := \{\theta \in \Theta^\delta : \widehat{m}(\theta) \leq 0\}.$$

In analogy to part 1, we establish the conclusion by the following steps:

$$\begin{aligned} \sqrt{\mathcal{W}_n} &= \sup_{\theta \in \Theta_0} \sqrt{nd}(\theta, \widehat{\Theta}_{0,\delta})/\widehat{w}(\theta) \stackrel{(1)}{=} \sup_{\theta \in \partial\Theta_0} [(\mathbb{G}_n(\theta) + \sqrt{nm}(\theta))/\widehat{w}(\theta) + o_p(1)]_+ \\ &\stackrel{(2)}{=} \sup_{\theta \in \partial\Theta_0} [(\mathbb{G}_n(\theta) + \sqrt{nm}(\theta))/w(\theta) + o_p(1)]_+ \\ &\rightsquigarrow (3) \sup_{\theta \in \partial\Theta_0} [(\mathbb{G}(\theta) + \sqrt{nm}(\theta))/w(\theta)]_+, \end{aligned}$$

where steps (1)-(3) are proven as follows:

To establish step (1), we first show that

$$\sup_{\theta \in \Theta_0} \sqrt{nd}(\theta, \widehat{\Theta}_{0,\delta})/\widehat{w}(\theta) = \sup_{\Theta_n} \sqrt{nd}(\theta, \widehat{\Theta}_{0,\delta})/\widehat{w}(\theta) \tag{A.4}$$

holds with probability approaching 1, where

$$\Theta_n = \{\theta \in \Theta_0 : d(\theta, \partial\Theta_0)/w(\theta) \leq \kappa_n\},$$

and κ_n is some sequence of positive random variables converging to zero in probability, $\kappa_n = o_p(1)$. Note that right hand side of (A.4) is less than or equal to the left hand side of (A.4) by construction, so we only need to show that w.p.a.1, the right hand side can not be less. To this end, fix some $\kappa > 0$ and note that using the same line of reasoning as for part 1, we have

$$\zeta(\kappa) := \sup\{m(\theta) : \theta \in \Theta_0, d(\theta, \partial\Theta_0) \geq \kappa\} < 0.$$

Furthermore, for every $\zeta < 0$ we have that with probability approaching 1, $\sup_{\theta \in \Theta} |\widehat{m}(\theta) - m(\theta)| < -\zeta$, and therefore each $\theta \in \Theta_0$ with $d(\theta, \partial\Theta_0) \geq \kappa$ is included in $\widehat{\Theta}_{0,\delta}$. In that event, we can only have $d(\theta, \widehat{\Theta}_{0,\delta}) > 0$ for values of θ within a distance κ of the boundary $\partial\Theta_0$. Since κ can be chosen arbitrarily small, we can choose a sequence $\kappa_n = o(1)$ such that the right-hand side in (A.4) holds with probability converging to one.

Next, let θ_n be a sequence such that

$$\sup_{\theta \in \Theta_0} \sqrt{n}d(\theta, \widehat{\Theta}_{0,\delta})/\widehat{w}(\theta) = \sqrt{n}d(\theta_n, \widehat{\Theta}_{0,\delta})/\widehat{w}(\theta_n).$$

Note that since $d(\theta, \widehat{\Theta}_{0,\delta})$ and $\widehat{w}(\theta)$ are continuous functions of θ , $\widehat{w}(\theta)$ is bounded away from zero with probability approaching 1, so that since Θ_0 is compact, the supremum is attained with probability approaching 1. Since $\delta > 0$ is fixed and by Lemma 3 with probability approaching 1, we have $d(\bar{\theta}_n, \theta_n)/\widehat{w}(\theta) < \delta$, so that $\bar{\theta}_n$ is contained in the interior of Θ^δ w.p.a.1.

It follows from Lemma 4 that

$$\sqrt{n}d(\theta_n, \widehat{\Theta}_{0,\delta}) = [\|\nabla_\theta m(\theta_n)\|^{-1}(\mathbb{G}_n(\theta_n) + \sqrt{n}m(\theta_n)) + o_p(1)]_+$$

Hence, dividing by $\widehat{w}(\theta)$ and using stochastic equicontinuity of $\mathbb{G}_n(\theta)$, we can follow the same line of reasoning as in part 1 to obtain that with probability approaching 1,

$$\sup_{\theta \in \Theta_0} \sqrt{n}d(\theta, \widehat{\Theta}_{0,\delta})/\widehat{w}(\theta) \leq \sup_{\theta \in \partial\Theta_0, \theta + \eta \in \Theta_0, \|\eta\| \leq \kappa_n} \left[\frac{\mathbb{G}_n(\theta) + \sqrt{n}m(\theta + \eta)}{\|\nabla_\theta m(\theta)\|\widehat{w}(\theta + \eta)} + o_p(1) \right]_+$$

Noting again that $\sqrt{n}m(\theta + \eta) \leq 0$ for $\theta + \eta \in \Theta_0$, we can bound the expression by

$$\sup_{\theta \in \Theta_0} \sqrt{n}d(\theta, \widehat{\Theta}_{0,\delta})/\widehat{w}(\theta) \leq \sup_{\theta \in \partial\Theta_0, \theta + \eta \in \Theta_0, \|\eta\| \leq \kappa_n} \left[\frac{\mathbb{G}_n(\theta)}{\|\nabla_\theta m(\theta)\|\widehat{w}(\theta + \eta)} + o_p(1) \right]_+$$

where again the $o_p(1)$ term is the same as in the first inequality.

Therefore, we can use uniform convergence of the weighting function \widehat{w} to w from Condition R.3 and the continuous mapping theorem to obtain

$$\mathcal{W}_n \rightsquigarrow \sup_{\theta \in \partial\Theta_0} [(\|\nabla_\theta m(\theta)\|w(\theta))^{-1} \mathbb{G}(\theta)]_+^2 \quad (\text{A.5})$$

where the intermediate steps are analogous to equalities (2) and (3) in the proof for \mathcal{L}_n .

PART 3. (Continuity of the Limit Distributions). The continuity of the distribution function \mathcal{L} on $(0, \infty)$ follows from Davydov et al. (1998) and from the assumption that

the covariance function of \mathbb{G} is non-degenerate, i.e. $\inf_{\theta \in \Theta} \text{Var}(\mathbb{G}(\theta)) > 0$. The probability that \mathcal{L} is greater than zero is equal to the probability that $\max_j \sup_{\theta \in \Theta} \mathbb{G}_j(\theta) > 0$, which is greater than the probability that $\mathbb{G}_{j'}(\theta') > 0$ for some fixed j' and θ' , but the latter is equal to $1/2$. Therefore the claim follows. The claim of continuity of the distribution function of \mathcal{W} on $(0, \infty)$ follows similarly. \square

A.2. Proof of Corollary 1. This corollary immediately follows from the assumed conditions and from the comments given in the main text preceding the statement of Corollary 1. \square

A.3. Proof of Theorem 2. We have that $\text{P}[\Theta_0 \subseteq R_{LR}] = \text{P}[\mathcal{L}_n \leq \widehat{k}(1 - \alpha)]$ by the construction of the confidence region. We then have that for any $\alpha < 1/2$ that $k(1 - \alpha)$ is a continuity point of the distribution function of \mathcal{L} , so that for any sufficiently small ϵ

$$\begin{aligned} \text{P}[\mathcal{L}_n \leq \widehat{k}(1 - \alpha)] &\leq \text{P}[\mathcal{L}_n \leq k(1 - \alpha) + \epsilon] + o(1) \rightarrow \text{P}[\mathcal{L} \leq k(1 - \alpha) + \epsilon], \\ \text{P}[\mathcal{L}_n \leq \widehat{k}(1 - \alpha)] &\geq \text{P}[\mathcal{L}_n \leq k(1 - \alpha) - \epsilon] - o(1) \rightarrow \text{P}[\mathcal{L} \leq k(1 - \alpha) - \epsilon]. \end{aligned}$$

Since we can set ϵ as small as we like and $k(1 - \alpha)$ is a continuity point of the distribution function of \mathcal{L} , we have that

$$\text{P}[\mathcal{L}_n \leq \widehat{k}(1 - \alpha)] \rightarrow \text{P}[\mathcal{L} \leq k(1 - \alpha)] = (1 - \alpha).$$

We can conclude similarly for the W -statistic \mathcal{W}_n . \square

A.4. Proof of Theorem 3. We will give the proof only for confidence regions based on the LR statistic. The arguments for the Wald-type confidence sets are completely analogous. Let δ_n be a null sequence where $\delta_n > 0$ for all n and $\sqrt{n}\delta_n \rightarrow \infty$. In order to show convergence with respect to Hausdorff distance, we establish that with probability approaching 1, (a) $R_{LR} \subset \Theta_0^{\delta_n}$, and (b) $\Theta_0 \subset R_{LR}^{\delta_n}$, where for $\delta > 0$, $A^\delta := \{x \in \mathbb{R}^k : d(x, A) < \delta\}$ denotes the δ -expansion of a set A in \mathbb{R}^k .

To prove statement (a), consider a sequence $\theta_n \in \Theta \setminus \Theta_0^{\delta_n}$. We have to show that $\theta_n \notin R_{LR}$ w.p.a.1 as n increases: For any fixed $\bar{\delta} > 0$, let $\eta := \inf_{\theta: d(\theta, \Theta_0) \geq \bar{\delta}} m(\theta)$. Since Θ is compact and $m(\theta)$ is continuous by R.2, it follows from the definition of Θ_0 that $\eta > 0$. Hence, along any sequence θ_n such that $d(\theta_n, \Theta_0) \geq \bar{\delta}$ for all n , we have $\sqrt{n}m(\theta_n) \rightarrow \infty$. Since $\theta \in R_{LR}$ if and only if for the critical value k $\mathcal{LR}_n \leq \widehat{k}$, where \widehat{k} is tight for any n large enough, it follows that $\theta_n \notin R_{LR}$ w.p.a.1. Hence we can restrict our attention to (sub-) sequences for which $d(\theta_n, \Theta_0) \rightarrow 0$.

Now suppose that $d(\theta_n, \Theta_0) \rightarrow 0$, and $\sqrt{n}d(\theta_n, \Theta_0) \rightarrow \infty$. Let θ_n^* be the projection of θ_n onto $\partial\Theta_0$, where θ_n^* need not converge to a particular point. Then using the same steps as in the proof of Lemma 4 it follows from a mean-value expansion of $m(\theta)$ around θ_n and continuity of the gradient $\nabla_{\theta}m(\theta)$ that $\sqrt{n}m(\theta_n) \rightarrow \infty$. Hence the LR-statistic diverges to infinity, and $\theta_n \notin R_{LR}$ w.p.a.1.

It remains to check claim (b), namely that $\Theta_0 \subset R_{LR}^{\delta_n}$. To this end, consider a sequence $\theta_n \in \Theta_0$, and let $q_n = \widehat{s}(\theta) \sqrt{\widehat{k}(1-\alpha)/n}$. Also define $\bar{\theta}_n \equiv \arg \min_{\widehat{m}(\theta) \leq q_n} \|\theta_n - \theta\|^2$, so that $d(\theta_n, \widehat{\Theta}_{0,\delta}) = \|\theta_n - \bar{\theta}_n\|$. Now note that Lemma 3 implies that $d(\bar{\theta}_n, \Theta_0) \xrightarrow{p} 0$, so that we can apply Lemma 4 to obtain

$$\sqrt{n}d(\bar{\theta}_n, \theta_n) = [\|\nabla_{\theta}m(\theta_n)\|^{-1}(\mathbb{G}_n(\theta_n) + \sqrt{n}(m(\theta_n) - q_n)) + o_p(1)]_+$$

By assumption R.2, the norm of $\nabla_{\theta}m(\theta_n)$ is bounded away from zero, and $\mathbb{G}_n(\theta_n)$ is stochastically bounded by Condition R.1. Since $\widehat{s}(\theta)$ and \widehat{k} are stochastically bounded by R.3, $\sqrt{n}q_n$ is also stochastically bounded.

Noting that $m(\theta_n) \leq 0$, it follows that $\sqrt{n}(\bar{\theta}_n - \theta_n) = O_p(1)$. Therefore

$$\sup_{\theta \in \Theta_0} \sqrt{n}d(\theta, R_{LR}) = O_p(1)$$

so that for any sequence δ_n such that $\sqrt{n}\delta_n \rightarrow \infty$, we have $d(\theta, R_{LR})/\delta_n \rightarrow 0$. Hence $\Theta_0 \subset R_{LR}^{\delta_n}$ w.p.a.1, which concludes the proof \square

A.5. Proof of Corollary 2. This corollary immediately follows from the assumed conditions and Corollary 1. \square

A.6. Proof of Theorem 4. In what follows let A denote an absolute positive constant. We have by definition of the Kantarovich-Rubenstein metric that

$$\mathbb{E}_{\mathcal{Q}_{V^*}}[\varphi(V^*)] - \mathbb{E}_{\mathcal{Q}_V}[\varphi(V)] = o_p(1) \text{ uniformly in } \varphi \in \text{BL}_1(C(\Theta)).$$

This implies that

$$\mathbb{E}_{\mathcal{Q}_{V^*}}[\varphi([V^*]_+^2)] - \mathbb{E}_{\mathcal{Q}_V}[\varphi([V]_+^2)] = o_p(1) \text{ uniformly in } \varphi \in \text{BL}_1(C(\Theta)),$$

since the composition $\varphi \circ [\cdot]_+^2 \in A \cdot \text{BL}_1(C(\Theta))$ for $\varphi \in \text{BL}_1(C(\Theta))$. This further implies that

$$\mathbb{E}_{\mathcal{Q}_{V^*}}[\varphi'(\sup_{R_n} [V^*]_+^2)] - \mathbb{E}_{\mathcal{Q}_V}[\varphi'(\sup_{R_n} [V]_+^2)] = o_p(1) \text{ uniformly in } \varphi' \in \text{BL}_1(\mathbb{R}),$$

since the composition $\varphi'(\sup_{R_n}[\cdot]_+^2) \in A \cdot \text{BL}_1(C(\Theta))$ for $\varphi' \in \text{BL}_1(\mathbb{R})$ and R_n denoting any sequence of closed non-empty subsets in Θ . Therefore, by the Extended Continuous Mapping Theorem,

$$\mathbb{E}_{\mathcal{Q}_{V^*}}[\varphi'(\sup_{\widehat{\partial\Theta_0}}[V^*]_+^2)] - \mathbb{E}_{\mathcal{Q}_V}[\varphi'(\sup_{\widehat{\partial\Theta_0}}[V]_+^2)] = o_p(1) \text{ uniformly in } \varphi' \in \text{BL}_1(\mathbb{R}).$$

(Note that here we compute expectations over V^* and V taking $\widehat{\partial\Theta_0}$ as given; note that our bootstrap method treats $\widehat{\partial\Theta_0}$ as fixed.) Also note that any sequence of sets R_n converging to a set R , we have that

$$\begin{aligned} & |\mathbb{E}_{\mathcal{Q}_V}[\varphi'(\sup_{R_n}[V]_+^2) - \varphi'(\sup_R[V]_+^2)]| \\ & \leq \mathbb{E}_{\mathcal{Q}_V}[|\sup_{R_n}[V]_+^2 - \sup_R[V]_+^2| \wedge 1] = o_p(1) \text{ uniformly in } \varphi' \in \text{BL}_1(\mathbb{R}), \end{aligned}$$

since

$$\sup_{R_n}[V]_+^2 - \sup_R[V]_+^2 = o_p(1)$$

by stochastic equicontinuity of the process V . Since by Condition R.1-R.3 and Theorem 3, $\widehat{\partial\Theta_0}$ converges to $\partial\Theta_0$ in the Hausdorff distance, we have by the Extended Continuous Mapping Theorem:

$$|\mathbb{E}_{\mathcal{Q}_V}[\varphi'(\sup_{\widehat{\partial\Theta_0}}[V]_+^2) - \varphi'(\sup_{\partial\Theta_0}[V]_+^2)]| = o_p(1) \text{ uniformly in } \varphi' \in \text{BL}_1(\mathbb{R}),$$

where \mathcal{Q}_V computes the expectation over V , treating $\widehat{\partial\Theta_0}$ as fixed.

Combining the steps above, we conclude by the triangle inequality that:

$$|\mathbb{E}_{\mathcal{Q}_{V^*}}[\varphi'(\sup_{\widehat{\partial\Theta_0}}[V^*]_+^2) - \varphi'(\sup_{\partial\Theta_0}[V]_+^2)]| = o_p(1) \text{ uniformly in } \varphi' \in \text{BL}_1(\mathbb{R}),$$

which is the same as

$$\rho_K(\mathcal{Q}_{S^*}, \mathcal{Q}_S) = o_p(1).$$

(Note that the bootstrap random variable \mathcal{S}^* is computed having fixed $\widehat{\partial\Theta_0}$.)

It is known that the convergence $\rho_K(\mathcal{Q}_{S_n}, \mathcal{Q}_S) = o(1)$, for any sequence of laws \mathcal{Q}_{S_n} of a sequence of random variables \mathcal{S}_n defined on probability space $(\Omega', \mathcal{F}', P_n)$ implies the convergence of the distribution function

$$\Pr_{\mathcal{Q}_{S_n}}[\mathcal{S}_n \leq s] = \Pr_{\mathcal{Q}_S}[\mathcal{S} \leq s] + o(1),$$

at each continuity point $(0, \infty)$ of the mapping $s \mapsto \Pr[\mathcal{S} \leq s]$ and also convergence of quantile functions

$$\inf\{s : \Pr_{\mathcal{Q}_{S_n}}[\mathcal{S}_n \leq s] \geq p\} = \inf\{s : \Pr_{\mathcal{Q}_S}[\mathcal{S} \leq s] \geq p\} + o(1)$$

at each continuity point p of the mapping $s \mapsto \inf\{s : \Pr_{\mathcal{Q}_S}[\mathcal{S} \leq s] \geq p\}$. Recall from Theorem 1 that the set of continuity points necessarily includes the region $(0, 1/2)$.

By the Extended Continuous Mapping Theorem (see e.g. Theorem 18.11 in van der Vaart (1998)) we conclude that since $\rho_K(\mathcal{Q}_{S^*}, \mathcal{Q}_S) = o_p(1)$, we obtain the convergence in probability of the distribution function

$$\Pr_{\mathcal{Q}_{S^*}}[\mathcal{S}^* \leq s] = \Pr_{\mathcal{Q}_S}[\mathcal{S} \leq s] + o_p(1),$$

at each continuity point $(0, \infty)$ of the mapping $s \mapsto \Pr[\mathcal{S} \leq s]$ and also convergence in probability of the quantile functions

$$\inf\{s : \Pr_{\mathcal{Q}_{S^*}}[\mathcal{S}^* \leq s] \geq p\} = \inf\{s : \Pr_{\mathcal{Q}_S}[\mathcal{S} \leq s] \geq p\} + o_p(1),$$

at each continuity point p of the mapping $s \mapsto \inf\{s : \Pr_{\mathcal{Q}_S}[\mathcal{S} \leq s] \geq p\}$. \square

A.7. Proof of Proposition 1. First note that γ_{HJ} and γ_{MF} are continuous, differentiable functions of the elements in v, Σ , and B , so that C.1 follows from a CLT for $\hat{v}, \text{vec}(\hat{\Sigma}), \text{vec}(\hat{B})$ and the delta-rule.

Next, we check Condition C.2 for the HJ bounds. Recall that the moment function defining the HJ bound was given by

$$m_{HJ}((\mu, \sigma)', \gamma_{HJ}) = \sqrt{S_{vv}\mu^2 - 2S_{v1}\mu + S_{11}} - \sigma$$

with S_{vv}, S_{v1}, S_{11} as defined in section 1.

Since the derivative of $m_{HJ}((\mu, \sigma)', \gamma_{HJ})$ with respect to σ is equal to minus one, the lower bound on the norm of the gradient holds for all $\theta \in \Theta$. Also, since the eigenvalues of Σ are bounded away from zero, and the elements of $|v|$ are bounded, $|S_{vv}|$ and $|S_{v1}|$ are also bounded. Furthermore, we can easily verify that

$$S_{vv}\mu^2 - 2S_{v1}\mu + S_{11} \geq \frac{S_{vv}S_{11} - S_{v1}^2}{S_{vv}}.$$

for all values of μ . Now, since the eigenvalues of Σ^{-1} are bounded away from zero, and $v'v1'_N1_N - (1'_Nv)^2$ is bounded from below by a positive constant by assumption, we can also bound $S_{vv}\mu^2 - 2S_{v1}\mu + S_{11}$ away from zero.

Noting that the parameter space $\Theta^\delta \times \Gamma$ is compact, it follows that the second derivatives of $m^{(HJ)}(\theta, \gamma_{HJ})$ with respect to γ_{HJ} and θ , respectively, are bounded. In particular, this implies Lipschitz continuity of the gradients $\nabla_\gamma m^{(HJ)}(\theta, \gamma_{HJ})$ and $\nabla_\theta m^{(HJ)}(\theta, \gamma_{HJ})$.

The arguments for the moment function $m_{MF}(\theta, \gamma_{MF})$ for the MF mean-variance set for asset portfolios are completely analogous. \square

A.8. Proof of Corollary 3. First note that by Proposition 1, Conditions C.1 and C.2 hold. Hence, in order to prove this corollary it suffices to show that

$$\rho_K(\mathcal{Q}_{\hat{v}Z^*}, \mathcal{Q}_{t'Z}; C(\Theta)) = o_p(1).$$

Without loss of generality we can take $\sup_\theta \|\hat{t}(\theta)\| \leq 1$ and $\sup_\theta \|t(\theta)\| \leq 1$. The claim will follow from

$$\rho_K(\mathcal{Q}_{\hat{v}Z^*}, \mathcal{Q}_{t'Z}; C(\Theta)) \leq \rho_K(\mathcal{Q}_{\hat{v}Z^*}, \mathcal{Q}_{\hat{v}Z}; C(\Theta)) + \rho_K(\mathcal{Q}_{\hat{v}Z}, \mathcal{Q}_{t'Z}; C(\Theta)) = o_p(1).$$

That $\rho_K(\mathcal{Q}_{\hat{v}Z^*}, \mathcal{Q}_{\hat{v}Z}; C(\Theta)) = o_p(1)$ follows immediately from $\rho_K(\mathcal{Q}_{Z^*}, \mathcal{Q}_Z) = o_p(1)$ and $\varphi(\hat{t} \cdot) \in \text{BL}_1(\mathbb{R}^k)$. Indeed,

$$\begin{aligned} |\varphi(\hat{t}(\theta)'z(\theta)) - \varphi(\hat{t}(\theta)'\bar{z}(\theta))| &\leq \sup_\theta |\hat{t}(\theta)'(z(\theta) - \bar{z}(\theta))| \wedge 2 \\ &\leq [(\sup_\theta \|\hat{t}(\theta)\| \sup_\theta \|z(\theta) - \bar{z}(\theta)\|) \wedge 2] \\ &\leq [\sup_\theta \|z(\theta) - \bar{z}(\theta)\| \wedge 2]. \end{aligned}$$

That $\rho_K(\mathcal{Q}_{\hat{v}Z}, \mathcal{Q}_{t'Z}; C(\Theta)) = o_p(1)$ follows because uniformly in $\varphi \in \text{BL}_1(C(\Theta))$

$$\begin{aligned} |\mathbb{E}_{\mathcal{Q}_Z}[\varphi(\hat{t}'Z)] - \varphi(t'Z)| &\leq \mathbb{E}_{\mathcal{Q}_Z}[\sup_\theta |(\hat{t}(\theta) - t(\theta))'Z(\theta)| \wedge 2] \\ &\leq \mathbb{E}_{\mathcal{Q}_Z}[\sup_\theta \|\hat{t}(\theta) - t(\theta)\| \sup_\theta \|Z(\theta)\| \wedge 2] = o_p(1), \end{aligned}$$

where $\mathbb{E}_{\mathcal{Q}_Z}$ computes the expectation over Z , treating \hat{t} as fixed. \square

A.9. Proof of Proposition 3. For the LR-type statistic, it is sufficient to notice that $\mathcal{L}_n(\theta; \eta) = \mathcal{L}_n(\eta^{-1}(\eta(\theta))) = \frac{\sqrt{n}\hat{m}_n(\theta)}{\hat{s}(\theta)}$, and therefore only depends on quantities evaluated at θ . Hence,

$$\mathcal{L}_n = \sup_{\theta \in \Theta_0} \mathcal{L}_n(\theta) = \sup_{\eta \in \eta(\Theta_0)} \mathcal{L}_n(\eta^{-1}(\eta)),$$

and $\phi_{\mathcal{L}_n}(\theta; \eta)$ is invariant with respect to parameter transformations. Asymptotic similarity follows from Condition R.1 and the continuous mapping theorem.

Next, define $\widehat{H}_{0,\delta} = \{\eta \in \eta(\Theta^\delta) : \widehat{m}(\eta) \leq 0\}$, noting that $\widehat{H}_{0,\delta} = \eta(\widehat{\Theta}_{0,\delta})$. We now distinguish three cases regarding the limit point of the sequence θ_n : (1) $\theta_0 \in \text{int } \Theta_0$, the interior of Θ_0 , (2) $\theta \in \Theta/\Theta_0$, and (3) $\theta \in \partial\Theta_0$.

Since the weighting functions $\widehat{w}(\theta; \eta)$ were assumed to satisfy Condition R.3, Theorem 3 implies that in case (1), $\phi_{W_n}(\theta; \eta) \rightarrow 0$ for all values of η , including the identity transformation, so that by continuity of $\mathcal{W}_n(\theta; \eta)$ in θ , we have

$$\begin{aligned} & \lim_{n \rightarrow \infty} \text{P}(\phi_{W_n}(\theta_n; \eta) \neq \phi_{W_n}(\theta_0; \eta)) \\ & \leq \max\left\{ \lim_{n \rightarrow \infty} \text{P}(\phi_{W_n}(\theta_0; \eta) = 1), \lim_{n \rightarrow \infty} \text{P}(\phi_{W_n}(\theta_0) = 1) \right\} = 0. \end{aligned}$$

Similarly, in case (2), we have

$$\begin{aligned} & \lim_{n \rightarrow \infty} \text{P}(\phi_{W_n}(\theta_n; \eta) \neq \phi_{W_n}(\theta_0; \eta)) \\ & \leq \max\left\{ \lim_{n \rightarrow \infty} \text{P}(\phi_{W_n}(\theta_0; \eta) = 0), \lim_{n \rightarrow \infty} \text{P}(\phi_{W_n}(\theta_0) = 0) \right\} = 0. \end{aligned}$$

Finally, consider the third case in which $\theta_0 \in \partial\Theta_0$: Using the expansion in Lemma 4, we obtain that

$$\begin{aligned} \mathcal{W}_n(\theta_n; \eta) &= \left[\frac{d(\eta(\theta_n), \widehat{H}_{0,\delta})}{w(\theta_n; \eta)} \right]^2 = \left[\frac{d(\eta(\theta_n), \eta(\widehat{\Theta}_{0,\delta}))}{w(\theta_n; \eta)} \right]^2 \\ &= \left[\frac{\mathbb{G}_n(\theta_n) + \sqrt{n}m(\theta_n)}{\|\nabla_\eta m_\eta(\theta_n; \eta)\| w(\theta_n; \eta)} + o_{\text{P}}(1) \right]_+^2. \end{aligned} \quad (\text{A.6})$$

Now suppose we split the sequence θ_n into three (possibly trivial) subsequences θ_{q_n} , θ_{r_n} , and θ_{s_n} , respectively, such that $\sqrt{n}m(\theta_{q_n}) = O(1)$, $\sqrt{n}m(\theta_{r_n}) \rightarrow \infty$, and $\sqrt{n}m(\theta_{s_n}) \rightarrow -\infty$. We can now analyze the behavior of $\phi_{W_n}(\cdot)$ separately along each of these subsequences. By Theorem 3,

$$\lim_{n \rightarrow \infty} \text{P}(\phi_{W_{r_n}}(\theta_{r_n}; \eta) = 0) = \lim_{n \rightarrow \infty} \text{P}(\phi_{W_{s_n}}(\theta_{s_n}; \eta) = 1) = 0$$

for all η , so that by the same arguments as before, the test is asymptotically invariant along these subsequences.

Finally, consider the sequence θ_{q_n} : if $\widehat{w}(\theta; \eta) = \frac{b(\theta)}{\|\nabla_\eta m(\eta(\theta))\|} + o_p(1)$, the expansion in (A.6) together with the continuous mapping theorem implies that along the subsequence

θ_{q_n} ,

$$\begin{aligned} \mathcal{W}_{q_n}(\theta_{q_n}; \eta) &= \left[\frac{\mathbb{G}_n(\theta_{q_n}) + \sqrt{q_n}m(\theta_{q_n})}{b(\theta_{q_n})} + o_p(1 + \sqrt{q_n}m(\theta_{q_n})) \right]_+^2 \\ &= \left[\frac{\mathbb{G}_n(\theta_{q_n}) + \sqrt{q_n}m(\theta_{q_n})}{b(\theta_{q_n})} + o_p(1) \right]_+^2. \end{aligned}$$

so that the leading term of this expression is a function of θ which does not depend on η .

Now note that for $\theta \in \partial\Theta_0$ we have $m(\theta) = 0$, so that by this expansion and stochastic equicontinuity of $\mathbb{G}(\theta)$ we have

$$\mathcal{W}_n(\eta) := \sup_{\theta \in \partial\Theta_0} \mathcal{W}_n(\theta; \eta) \rightsquigarrow \mathcal{W} := \sup_{\theta \in \partial\Theta_0} \left[\frac{\mathbb{G}(\theta)}{b(\theta)} \right]_+^2$$

where the limit does not depend on η . It follows that $\widehat{k}_{W_\eta}(1 - \alpha) = \widehat{k}_W(1 - \alpha)$ for all η , so that $\phi_{W_n}(\theta; \eta)$ is asymptotically invariant along θ_{q_n} with respect to transformations $\eta \in H^*$, which establishes the second conclusion.

To establish the last claim, consider the Wald statistic in (2.4) with weighting function $\widehat{w}(\theta) = \frac{s(\theta)}{\|\nabla_\theta m(\theta)\|} + o_p(1)$. By Condition R.1 and the approximation in (A.5), $\mathcal{W}_n(\theta)$ converges in distribution to $\max\{0, Z\}^2$, where $Z \sim N(0, 1)$, for all values of $\theta \in \partial\Theta_0$ and is therefore asymptotically pivotal on the boundary of Θ_0 . Hence for any fixed critical level k , the limit $\lim_{n \rightarrow \infty} P(\mathcal{W}_n(\theta) > k)$ is constant across all values of $\theta \in \partial\Theta_0$ for any value k .

REFERENCES

- Beresteanu, A. and Molinari, F. (2008). Asymptotic properties for a class of partially identified models. *Econometrica*.
- Bohrer, R. (1973). An optimality property of Scheffé bounds. *Annals of Statistics*.
- Britten-Jones, M. (1999). The sampling error in estimates of mean-variance efficient portfolio weights. *Journal of Finance*.
- Casella, G. and Strawderman, W. (1980). Confidence bands for linear regression with restricted predictor variables. *JASA*.
- Chernozhukov, V., Hong, H., and Tamer, E. (2007). Estimation and confidence regions for parameter sets in econometric models. *Econometrica*.
- Chetty, R. (2012). Bounds on elasticities with optimization frictions: A synthesis of micro and macro evidence on labor supply. *Econometrica*.

- Cochrane, J. H. (2005). *Asset Pricing*. Princeton University Press.
- Critchley, F., Marriott, P., and Salmon, M. (1996). On the differential geometry of the wald test with nonlinear restrictions. *Econometrica*.
- Davydov, Y., Lifshits, M., and Smorodina, N. (1998). *Local Properties of Distributions of Stochastic Functionals*. American Mathematical Society, Providence, RI.
- Fama, E. (1996). Multifactor portfolio efficiency and multifactor asset pricing. *The Journal of Financial and Quantitative Analysis*.
- Gibbons, M., Ross, S., and Shanken, J. (1989). A test on the efficiency of a given portfolio. *Econometrica*.
- Hansen, L. P. and Jagannathan, R. (1991). Implications of security market data for models of dynamic economies. *The Journal of Political Economy*.
- Hansen, L. P. and Singleton, K. (1982). Generalized instrumental variables estimation of nonlinear rational expectations models. *Econometrica*.
- Kaido, H. and Santos, A. (2011). Asymptotically efficient estimation of models defined by convex moment inequalities. working paper, BU and UCSD.
- Lehmann, E. and Romano, J. (2005). *Testing Statistical Hypotheses*. Springer.
- Lettau, M. and Ludvigson, S. (2009). Euler equation errors. *Review of Economic Dynamics*.
- Ludvigson, S. (2012). Advances in consumption-based asset pricing: Empirical tests. *in: George M. Constantinides and Milton Harris and Rene M. Stulz (eds.), Handbook of the Economics of Finance, vol. 2*.
- Markowitz, H. (1952). Portfolio selection. *The Journal of Finance*.
- Molchanov, I. S. (1998). A limit theorem for solutions of inequalities. *Scandinavian Journal of Statistics*.
- Naiman, D. (1984). Optimal simultaneous confidence bounds. *Annals of Statistics*.
- Peñaranda, F. and Sentana, E. (2010). Spanning tests in return and stochastic discount factor mean-variance frontiers: A unifying approach. working paper, UPF and CEMFI.
- Politis, D. N. and Romano, J. P. (1994). Large sample confidence regions based on subsamples under minimal assumptions. *The Annals of Statistics*.
- Scheffé, H. (1953). A method for judging all contrasts in the analysis of variance. *Biometrika*.
- Sentana, E. (2009). The econometrics of mean-variance efficiency tests: a survey. *Econometrics Journal*.

van der Vaart, A. W. (1998). *Asymptotic Statistics*. Cambridge University Press.

van der Vaart, A. W. and Wellner, J. A. (1996). *Weak Convergence and Empirical Processes*. Springer-Verlag New York.

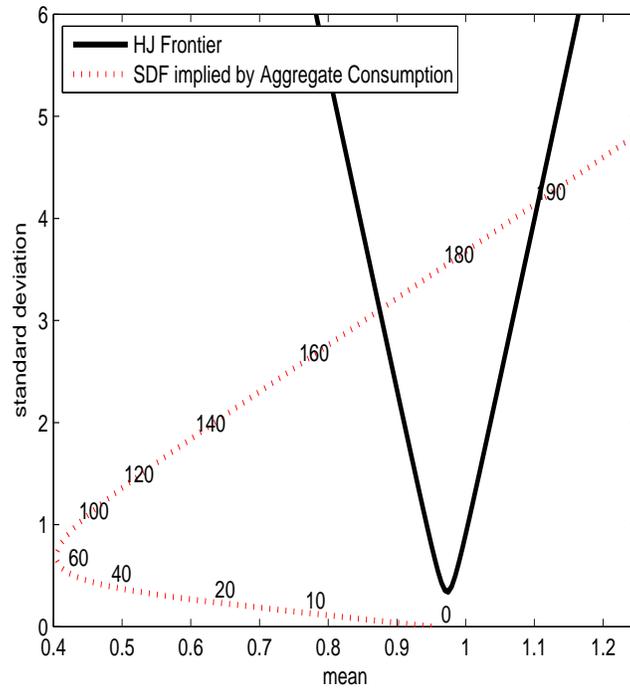


FIGURE 1. Estimated HJ Bounds with Stochastic Discount Factors implied by CRRA preferences, data labels represent values for the IES ρ

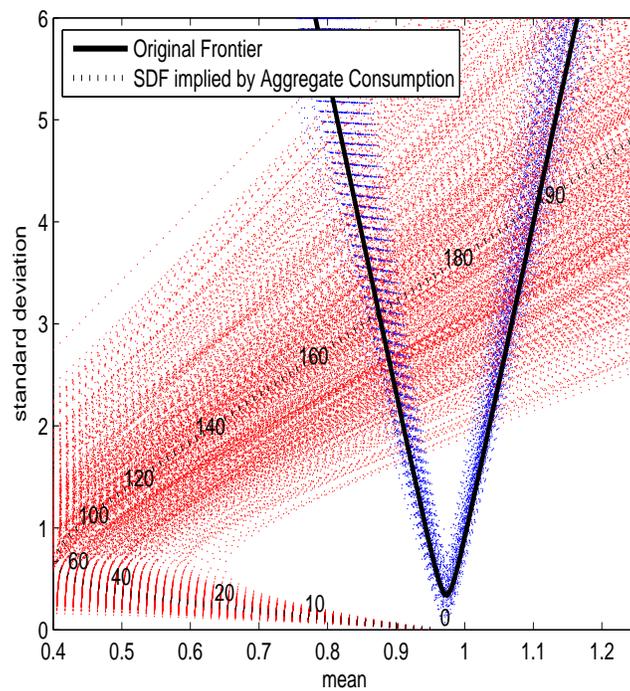


FIGURE 2. Estimated HJ Bounds, Mean-Standard deviation pairs based on Consumption SDF, and Bootstrap Draws

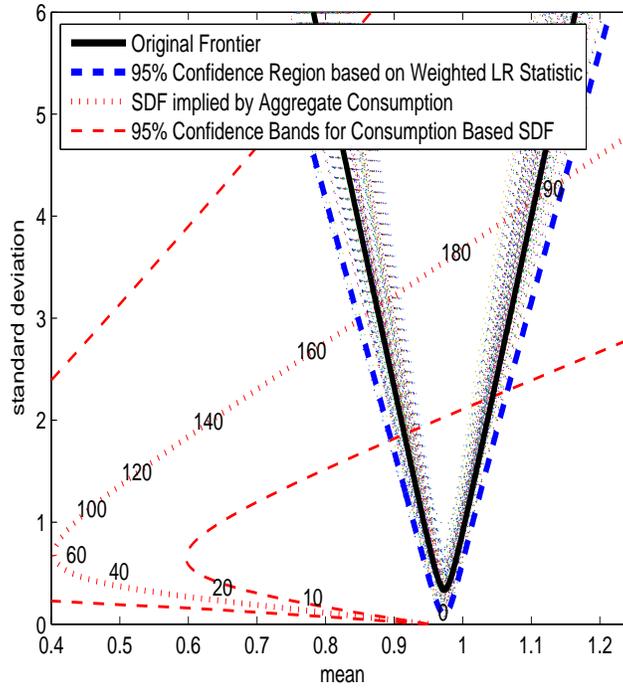


FIGURE 3. 95% Confidence Set for the HJ Set using LR Statistic. The 95% confidence set for mean and standard deviation of the consumption-based SDF was also constructed using the LR Statistic for the moment equality defining the corresponding mean-variance pairs.

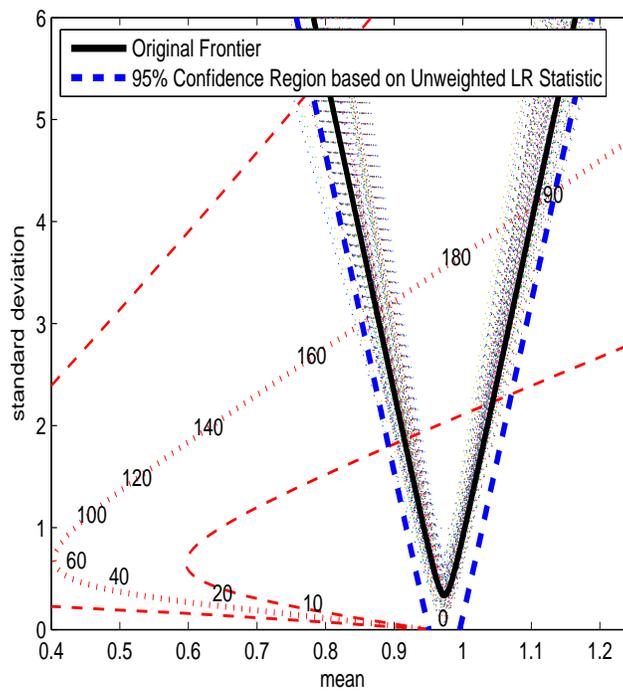


FIGURE 4. 95% Confidence Set for the HJ Set using Unweighted LR Statistic

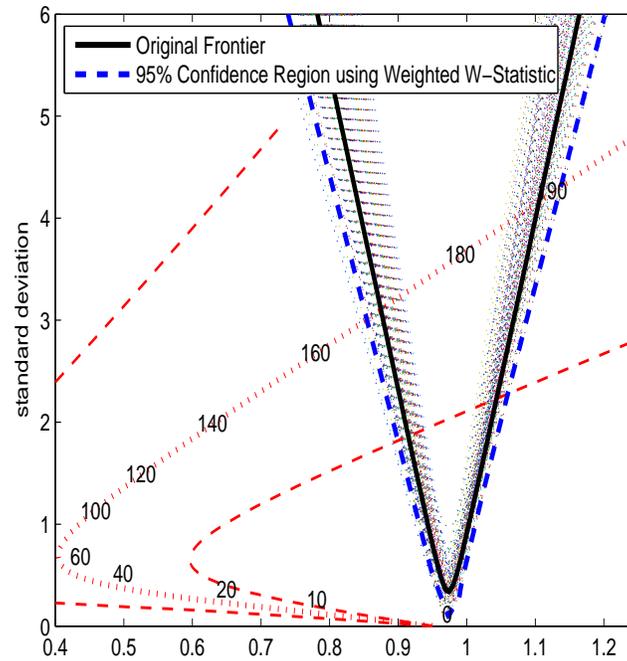


FIGURE 5. 95% Confidence Set for the HJ Set using Weighted W Statistic

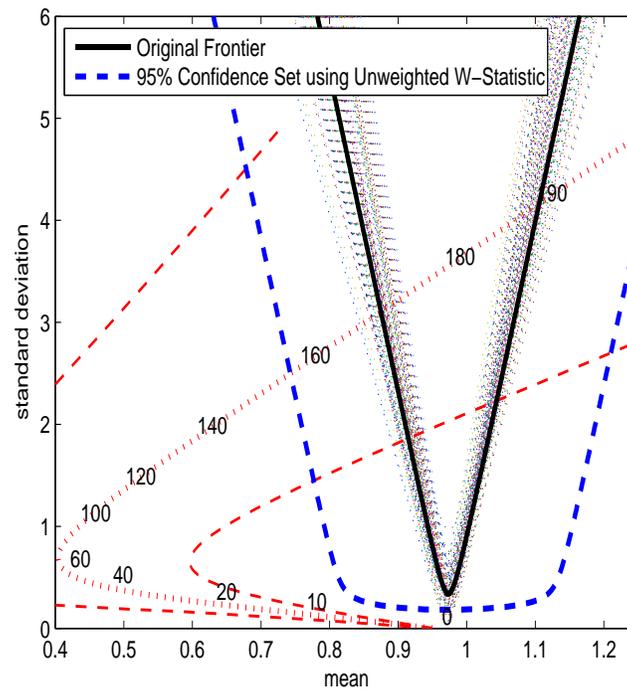


FIGURE 6. 95% Confidence Set for the HJ Set using Unweighted W Statistic (directed Hausdorff Distance)

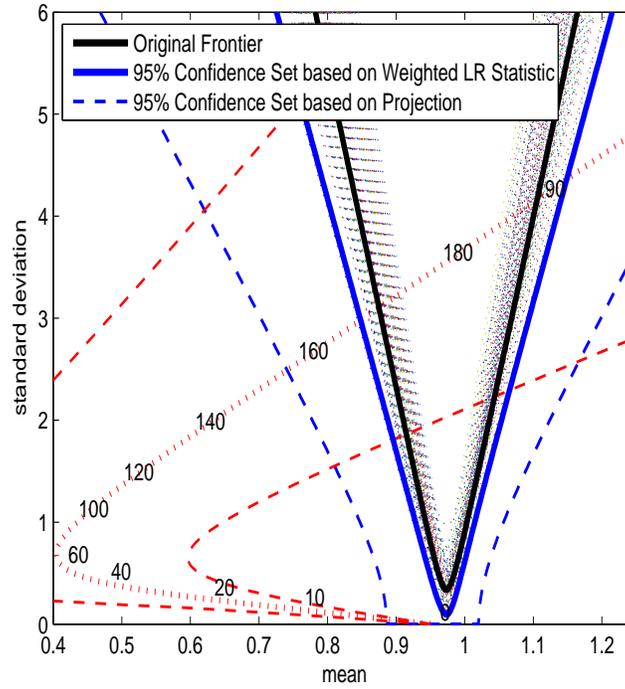


FIGURE 7. 95% Confidence Sets for the HJ Set using Weighted LR Statistic and Projection Approach

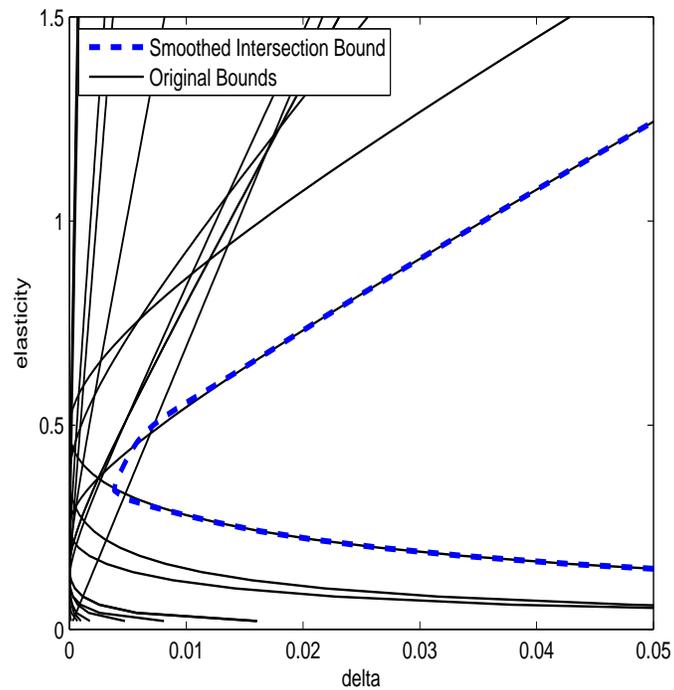


FIGURE 8. Bounds on Structural Labor Supply Elasticity based on Estimates from Different Studies, and their Smoothed Intersection Bound

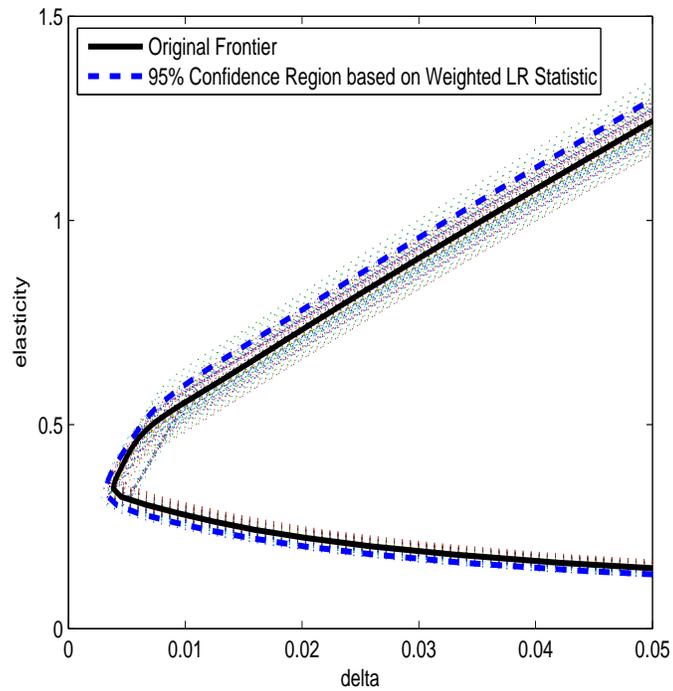


FIGURE 9. Smoothed Bound with 95% Confidence Set for the plausible pairs of the structural elasticity ε and friction δ using the LR Statistic

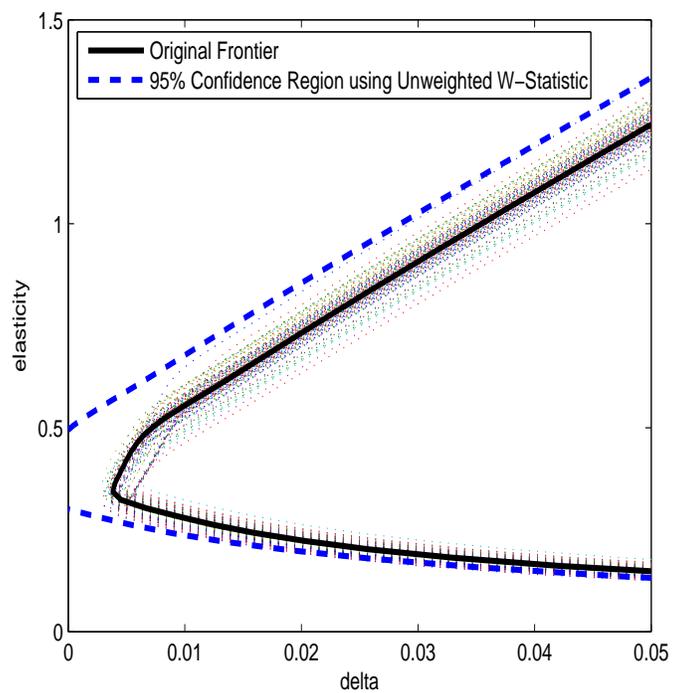


FIGURE 10. Smoothed Bound with 95% Confidence Set for the plausible pairs of the structural elasticity ε and friction δ using Unweighted W-Statistic