Testing Information Ordering for Strategic Agents

Sukjin Han, Hiroaki Kaido, Lorenzo Magnolfi

U of Bristol, Boston U, U of Wisconsin-Madison

June 2025

Workshop in honour of
Prof. Oliver Linton and Prof. Yoon-Jae Whang

Motivation

Many economic interactions are strategic

Researchers often bring models of games to data

- to estimate primitives and perform counterfactual simulations

One such primitive is the information available to players as they interact and generate the data

(i) information is needed to evaluate counterfactual policies, or

(ii) information may be of independent economic interest

- e.g., do politically connected firms get preferential info in procurement auctions? (Baltrunaite 20)

In either case, information structure prevailing in strategic interaction is seldom known to the researcher

An Example

Discrete game:

- $y_i \in \{1, 0\}$: "enter" or "not enter"
- firm i's profit: $\pi_i(y, \varepsilon_i; x, \theta)$
  - e.g. $\pi_i(y, \varepsilon_i; x, \theta) = x'\beta + \Delta y_{-i} + \varepsilon_i$
  - the payoff states ε_i are unobservable to the researcher

What do the players know about $\varepsilon \equiv (\varepsilon_1, \varepsilon_2)$?

- some players may know more than others
- they may know something in common

Predictions change depending on how the analyst specifies the information structure

With data and background alone, specifying info structure is hard

What This Paper Does

We represent information structures as high-dimensional nonparametric objects

We formulate statistical hypotheses to test whether the information structure prevailing in the data exceeds a certain baseline

We adopt the ordering of information structures, which respects incentive ordering (Bergemann & Morris 16)

We construct a bootstrap-based test that is asymptotically valid

We obtain as a by-product confidence set on payoff parameters

Application: we investigate information asymmetries in airline entry due to hubbing

Literature

Information ordering & Bayes correlated equilibria: Blackwell 53, 54, Kamenica & Gentzow 11, Bergemann & Morris 16;

Inference with general information structures: Grieco 14, Syrgkanis et al 19, Magnolfi & Roncoroni 23, Gualdani & Sinha 19;

Counterfactual predictions: Canen & Song 22, Bergemann et al 19;

Econometric tools (moment inequalities): Beresteanu & Molinari 08, Andrews & Soares 10, Beresteanu et al 11, Andrews & Barwick 12, Bontemps et al 12, Chernozhukov et al 13, Kaido & Santos 14, Bugni et al 15

Primitives

Setup

Primitives of a game:

- $i \in N$: players
- $y_i \in Y_i$: player i's action
- $x \in X$: game characteristics
- $\varepsilon_i \in \mathcal{E}_i$: payoff state; $\varepsilon \equiv (\varepsilon_i)_{i \in N} \sim F(\cdot; \theta)$: prior belief
- $\pi_i(y, \varepsilon_i; x, \theta)$: player i's payoff

What an analyst sees:

- $i \in N$: players
- $y_i \in Y_i$: player i's action
- $x \in X$: game characteristics

We focus on discrete games (i.e., Y_i is finite)

The players have common knowledge of the game, know (x, θ), but their knowledge of ε may be limited

Information Structure

Player i receives a private signal τ_i^x

$$\tau^x \equiv (\tau_1^x, \ldots, \tau_{|N|}^x) \sim P_{\tau|\varepsilon}^x$$

- τ_i^x is higher-order belief that carries info on payoff states ε

An information structure is a map from x to the conditional laws of the signals:

$$S : x \mapsto \left(\mathcal{T}^x, \left\{P_{\tau|\varepsilon}^x : \varepsilon \in \mathcal{E}\right\}\right)$$

We view $S(\cdot)$ as a nonparametric object

The payoff primitives and information structure define a game:

$$\Gamma^x(\theta, S)$$

Equilibrium Concept

The individuals play a Bayes Nash equilibrium (BNE):

- $\sigma_i : \mathcal{T}_i^x \to \Delta(Y_i)$
- A strategy profile $\sigma \equiv \left(\sigma_1, \ldots, \sigma_{|N|}\right)$ is a BNE of $\Gamma^x(\theta, S)$ if σ_i is a best response to σ_{-i} for all i

Define the set of BNE predictions:

$$Q_{\theta,S}{}^{BNE}(x) \equiv \left\{ q(\cdot|x) \in \Delta^{|Y|} \middle| q\left(y|x\right) = E[\sigma(y|\varepsilon,\tau)|x], \sigma \in BNE^x(\theta, S) \right\}$$

- the set of conditional choice probabilities (CCPs) induced by equilibria in $\Gamma^x(\theta, S)$
- $Q_{\theta,S}{}^{BNE}(x)$ requires the knowledge of S
- calculating this set requires finding all fixed points of best-response conditions

Information Ordering

Baseline Information Structure

Specifying S exactly can be hard

Instead, we consider testing if S is at least as informative as certain baseline S^r

Example. incomplete information: S_{Inc}

- τ_i reveals ε_i only and is not informative about ε_{-i}

Example. public signals: S_{Pub}

- $\varepsilon_i = \nu_i + \epsilon_i$
- for each player, τ_i reveals the opponent's shock ν_{-i}

Example. privileged signals: $S_{Priv,1}$

- τ_1 fully reveals $\varepsilon \equiv (\varepsilon_1, \varepsilon_2)$; τ_2 is only informative on ε_2

Example. complete information: S_{Comp}

- τ_i fully reveals $\varepsilon \equiv (\varepsilon_1, \varepsilon_2)$

Baseline Information

The information structure can vary across x (e.g., markets)

Example. privileged signals at hub airports

$$S^r(x) = \begin{cases} S_{Priv,1} & x_1 = 1 \\ S_{Inc} & x_1 = 0 \end{cases}$$

where $x_1 = 1\{\text{Airport} = \text{Hub for Player 1}\}$

Information Ordering

We want to state that the actual information structure S is at least as informative as a baseline S^r

- requires appropriate notion of informativeness
- Bergemann & Morris 16
- multi-agent generalization of Blackwell 51's information ordering

Information Ordering

Definition (Individual Sufficiency)

$S^1(x)$ is individually sufficient for $S^2(x)$ if there exist a combined information structure $S^*(x)$ such that

$$\tau_i^2|\tau_i^1, \tau_{-i}^1, \varepsilon \sim \tau_i^2|\tau_i^1, \ \forall i$$

That is, τ_i^2 is independent of $(\tau_{-i}^1, \varepsilon_i, \varepsilon_{-i})$ given τ_i^1.

▸ Combination

- $S^2(x)$ conveys no new information to any player about the payoff state and higher-order beliefs about the state

Write

$$S^1 \succeq S^2$$

if $S^1(x)$ is individually sufficient for $S^2(x)$ for all $x \in X$

Hypothesis Tests

Hypothesis

We test

$$H_0 : S \succeq S^r \quad \text{v.s.} \quad H_1 : S \not\succeq S^r$$

- H_0: players observe more info than S^r
- e.g., $S^r(x) = S_{Priv,1}(x)$ at hub airports

Testing Hypothesis

Recall, we assume CCPs are generated from BNE with S

How to contrast the data (captured in CCPs) with the hypothesis on information ordering?

To this end, we consider a solution concept (Bayes correlated equilibrium, Bergemann & Morris 16) that...

- incorporates the information ordering,
- which corresponds to the incentive ordering,
- and thus the ordering of equilibrium predictions

Bayes Correlated Equilibrium (BCE)

A Bayes correlated equilibrium (BCE) ν^x for the game $\Gamma^x(\theta, S^r)$ is a probability measure ν^x over actions profiles, payoff types, and signals that are:

1. consistent: for any measurable $A \subset \mathcal{E} \times \mathcal{T}$,

$$\int_A \int_{\mathcal{Y}} \nu^x(dy, d\varepsilon, dt) = \int_A P^x_{\tau|\varepsilon}(dt|\varepsilon) F(d\varepsilon; \theta_\varepsilon)$$

2. incentive compatible: for $y_i, \varepsilon_i, \tau_i$ s.t. $\nu^x (y_i \mid \varepsilon_i, \tau_i) > 0$,

$$E_{\nu^x} \left[\pi_i \left(y_i, y_{-i}, \varepsilon_i; x, \theta_\pi\right) \mid y_i, \tau_i\right] \geq E_{\nu^x} \left[\pi_i \left(y_i', y_{-i}, \varepsilon_i; x, \theta_\pi\right) \mid y_i, \tau_i\right] \quad \forall y_i' \in \mathcal{Y}_i$$

where $E_{\nu^x}[\cdot]$ is taken w.r.t. the conditional equilibrium distribution $\nu^x \left(y_{-i}, \varepsilon_{-i}, \tau_{-i} \mid y_i, \varepsilon_i, \tau_i\right)$

How Do We Use BCE

Here is how we interpret BCE

The individuals play a BNE under unknown info structure S

From the analyst's point of view, their behavior is consistent with the following story:

1. there's a baseline info structure S^r; the players may know more than S^r
2. a mediator observes $\varepsilon \sim F(\cdot; \theta_\pi)$ and $\tau \sim P_{\tau|\varepsilon}$ under S^r
3. the mediator draws $y \sim \nu(y|\tau, \varepsilon)$ and privately tells each i to play y_i
4. the players obey the mediator's recommendation

This view is convenient because we do not need to know the precise form of S (as long as $S \succeq S^r$)

Predictions

For a game $\Gamma^x(\theta, S)$, the set of BCE predictions is:

$$Q^{BCE}_{\theta,S}(x) \equiv \left\{ q(\cdot|x) \in \Delta^{|\mathcal{Y}|} \middle| q(y|x) = \int_{\mathcal{E}\times\mathcal{T}} \nu^x (y, d\varepsilon, d\tau) \,, \right.$$
$$\left. \nu^x \in BCE^x(\theta, S) \right\}$$

What's useful for us is the relationship between the BNE and BCE predictions:

Proposition

Suppose the data are generated by a BNE in $\Gamma^x(\theta, S)$ and $S \succeq S^r$
Then, for all $\theta \in \Theta$ and $x \in X$,

$$\underbrace{P_{y|x}}_{\text{identified}} \in Q_{\theta,S}{}^{BNE}(x) \subseteq \underbrace{Q^{BCE}_{\theta,S^r}(x)}_{\text{specified, convex}}$$

BCE Predictions

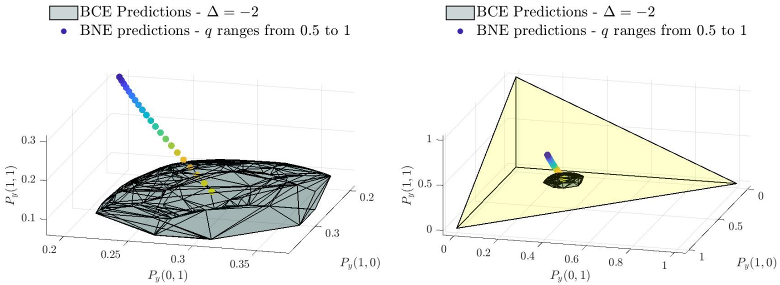

Figure: Colored dots: BNE CCPs with varying signal precisions ($S = S_{Inc}$ if $q = 0.5$; $S = S_{Pub}$ if $q = 1$)

Ordering of Information into Ordering of Functions

$$H_0 : S \succeq S^r \quad \text{v.s.} \quad H_1 : S \not\succeq S^r$$

If data are generated from a BNE under info structure S, then

$$P_{y|x} \in Q_{\theta,S}{}^{BNE}(x) \subseteq Q_{\theta,S^r}{}^{BCE}(x), \ \forall x \in X \tag{1}$$

- can detect the violation of H_0 if observed CCP is outside BCE prediction

▸ Testability

(1) can be translated into an ordering of functions

$$b' P_{y|x} \leq h(b, Q_{\theta,S^r}{}^{BCE}(x)), \ \forall b \in \mathbb{B}_x \text{ and } \forall x \in X$$

where $h(\cdot, A)$ is the support function of set A

▸ Support Function

- this allows us to work with functions rather than sets

Test Statistic

Let $(y^n, x^n) \equiv (y_\ell, x_\ell)_{\ell=1}^n$ be random sample drawn across markets

Let $\hat{P}_{n,x}$ be a vector of empirical CCPs

Define

$$T_n(\theta) \equiv \sup_{x \in X} \sup_{b \in \mathbb{B}_{n,x}} \sqrt{n}\{b' \hat{P}_{n,x} - h(b, Q_{\theta,S^r}{}^{BCE}(x))\}$$

where $\mathbb{B}_{n,x}$ is a "unit ball" with $\widehat{\text{AsyVar}}(\hat{P}_{n,x})$-weighted norm

- $T_n(\theta) = 0$ if $\hat{P}_{n,x} \in Q_{\theta,S^r}{}^{BCE}(x)$ and $T_n(\theta) > 0$ otherwise
- using the variance-weighted ellipsoid $\mathbb{B}_{n,x}$ has the effect of studentization
- easy to compute via convex quadratic program

Computation

Bootstrap

Consider a (empirical) bootstrap version of $T_n(\theta)$

$$T_n^*(\theta) \equiv \sup_{(b,x)\in\Psi_{n,\theta}} \{\mathbb{G}_n^*(b,x)\}$$

where

- $\mathbb{G}_n^*(b,x) \equiv \sqrt{n} b'(\hat{P}_{n,x}^* - \hat{P}_{n,x})$: bootstrapped empirical process
- $\Psi_{n,\theta} \equiv \{(b,x) : b'\hat{P}_{n,x} - h(b, Q_{\theta,S^r}^{BCE}(x)) \geq -\tau_n\}$: a conservative estimator of the "contact set" Ψ_θ

$$\Psi_\theta \equiv \{(b,x) : b' P_{y|x} = h(b, Q_{\theta,S^r}^{BCE}(x))\}$$

- Andrews & Soares 10; Chernozhukov et al 13

Define the bootstrap p-value by

$$p_n(\theta) \equiv P^*(T_n^*(\theta) > T_n(\theta)|y^n, x^n)$$

Asymptotic Size Control

Let $\phi(y^n, x^n) \equiv 1\{\sup_{\theta\in\Theta} p_n(\theta) \leq \alpha\}$

Theorem (Asymptotic Size)

Under H_0,

$$\limsup_{n\to\infty} \sup_{P\in\mathcal{P}} E_P[\phi] \leq \alpha$$

where $\mathcal{P}$ is the set of distributions of (y, x) that satisfy our assumptions and regularity conditions.

Extension: Testing Multiple Hypotheses

The analyst may consider testing a sequence of hypotheses

- to refine her understanding of the game's info structure

Example. Suppose $H_0 : S \succeq S_{Priv}$

- even if H_0 is not rejected, it is not definitive that $S = S_{Priv}$
- it is plausible that $S = S_{Comp}$, because $S_{Comp} \succeq S_{Priv}$

Instead, suppose the analyst test two nulls of the form:

$$H_{0,j} : S \succeq S_j^r \text{ for } j = 1, 2$$

where $S_1^r = S_{Comp}$ and $S_2^r = S_{Priv}$

- suppose $H_{0,2}$ is not rejected (as before) while $H_{0,1}$ is
- then, stronger evidence towards the player's privileged info

We introduce a modified version of Holm 79

- to control for family-wise error rate (and thus asymptotic size)

Monte Carlo Experiments

Simulation Design

Two-player binary response game:

$$\pi_i(y, \varepsilon_i; x, \theta_\pi) = x\beta + \Delta y_{-i} + \varepsilon_i, \quad i = 1, 2$$

- $\varepsilon_i = \nu_i + \epsilon_i$ with $\nu_i \in \{-\eta, \eta\}$; $x \in \{-M, M\}$

Under S, each player receives a signal about ν_{-i}

$$q \equiv P(t_i = \bar{\nu}_{-i} | \nu_{-i} = \bar{\nu}_{-i})$$

The precision of the signal is controlled by q

- $q \to \frac{1}{2}$: uninformative signal (i.e., S_{Inc})
- $q \to 1$: signal reveals ν_{-i} (i.e., S_{Pub})

We select a BNE and generate a sample of size $n = 1000$

We test $H_0 : S \succeq S_{Pub}$ against $H_1 : S \not\succeq S_{Pub}$

BCE Predictions

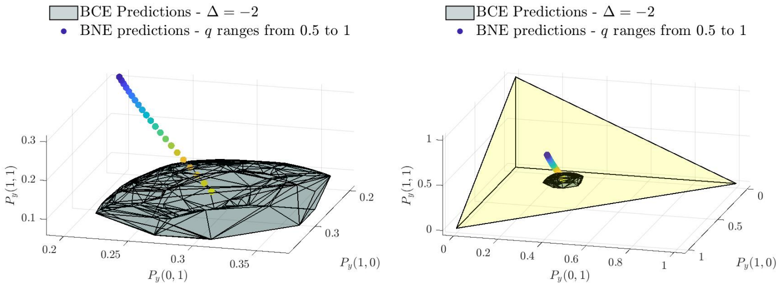

Figure: Colored dots: BNE CCPs with varying signal precisions ($S = S_{Inc}$ if $q = 0.5$; $S = S_{Pub}$ if $q = 1$)

Power Properties

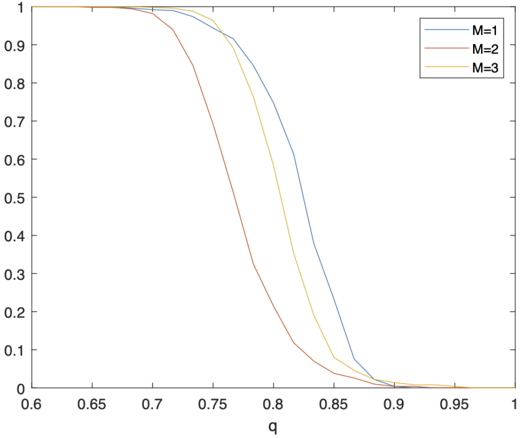

Figure: The Rejection Probability of the Test

An Empirical Application

An Empirical Question

Consider hubbing in the airline industry in the US

Q: When considering entry on potential new routes that originate in its hub, does the hub airline benefit from the superior ability to forecast demand and a better understanding of costs?

Data:

- Department of Transportation's Origin and Destination Survey (DB1B) and Domestic Segment (T-100) database
- markets (cross-sectional units): origin and destination airports in a given quarter
- potential entrants: American (AA), Delta (DL), United (UA), Southwest (WN), a medium-size airline, and a low-cost carrier

Players

We focus on hubs for AA, DL, UA, WN

We aggregate airlines into three players for each market (e.g., Atlanta - Airport X)

- hub airline (e.g., DL)
- non-hub airline (e.g., AA, UA, WN)
- non-major airline (e.g., midsize, LCC)

Observable covariates include airport presence (Berry 92), cost (Ciliberto & Tamer 09), and market characteristics (population, per capita income)

Hypothesis

We test aggregate null hypothesis

$$H_0 : S \succeq S_{Priv,Hub} \quad \text{v.s.} \quad H_1 : S \not\succeq S_{Priv,Hub}$$

and market-specific null hypothesis

$$H_{0,x} : S(x) \succeq S_{Priv,Hub}(x) \quad \text{v.s.} \quad H_{1,x} : S(x) \not\succeq S_{Priv,Hub}(x)$$

- the baseline information structure is $S_{Priv,Hub}$:
  - τ_{Hub} reveals ε;
  - τ_i for other i's only reveals their own payoff state
  - i.e., $\varepsilon_i = \nu_i + \epsilon_i$, and τ_{Hub} reveals ν_{-Hub}
- covariates x: market size, each player's market presence
  - total 16 market types

Informational Priviledge of Hub Airline

H_0 is rejected with $\inf_\theta\{T_n(\theta) - c^*_{0.05}(\theta)\} = 86.76$

- Bayesian optimization algorithm for $\inf_\theta$

$H_{0,x}$ is rejected for some (but not all) x

- "not rejected" even when hub airline has low market presence
- rejected in markets where hub and non-major airlines have high market presence

CS for Null Markets

Market →	0001	0010	0011	0111	1000	1001	1011	1101	$T - c$
Iter. 117	5.91	5.29	0.00	76.85	0.00	0.32	7.30	91.59	82.68
Iter. 122	5.91	6.86	0.00	80.01	0.76	3.89	9.17	94.68	85.90
Iter. 132	5.91	3.78	0.00	88.95	5.01	3.89	13.42	103.44	94.37
Iter. 135	1.30	4.28	0.00	89.67	6.75	3.88	15.16	104.15	94.93
Iter. 149	0.00	3.91	0.00	87.94	4.75	2.88	13.16	102.45	93.46

Table: Market-Specific Test Results (some columns omitted)

Conclusions

Concluding Remarks

The actual information structure of a strategic environment is a complex parameter

Nonetheless, it plays a crucial role in evaluating the model's empirical contents and making counterfactual predictions

The paper develops a test of hypotheses on the information structure

It will allow the researcher to

- investigate the players' information asymmetries;
- investigate how the info structure varies with market/game characteristics;
- use $Q^{BCE}_{\tilde{\theta},S^r}(\tilde{x})$ for counterfactual predictions

Thank You! ☺

Combining Signals

Definition (Combination)

The information structure (at x)

$$S^*(x) \equiv \left(\mathcal{T}^*_x, \left\{P^{*,x}_{\tau|\varepsilon} : \varepsilon \in \mathcal{E}\right\}\right)$$

is a combination of $S^1(x)$ and $S^2(x)$ if

$$\mathcal{T}^*_{i,x} = \mathcal{T}^1_{i,x} \times \mathcal{T}^2_{i,x}, \text{ for each } i$$

$$\int P^{*,x}_{\tau^*|\varepsilon}(\tau^1, \tau^2|\varepsilon)d\tau^1 = P^{2,x}_{\tau^2|\varepsilon}(\tau^2|\varepsilon) \text{ for each } \tau^2 \text{and } i$$

$$\int P^{*,x}_{\tau^*|\varepsilon}(\tau^1, \tau^2|\varepsilon)d\tau^2 = P^{1,x}_{\tau^1|\varepsilon}(\tau^1|\varepsilon), \text{ for each } \tau^1 \text{and } i$$

- we consider a coupling of the signals (given ε)

Return

Testability

Consider a simple example with two players:

- $\pi_i(y, \varepsilon_i) = y_i(-\Delta_j y_{-i} + \varepsilon_i)$ for $(\Delta_1, \Delta_2) \in \Theta = (0,1]^2$ and $\varepsilon_i \overset{iid}{\sim} U[-1,1]$

Suppose $S^r = S_C$ as the baseline

BCE predicts the sharp LB for the prob of e.g. $y = (1,0)$ as

$$LB_\Delta = \frac{1}{4}(1 + \Delta_2(1 - \Delta_1)) \geq 0.25, \quad \forall(\Delta_1, \Delta_2) \in \Theta$$

Let $\Delta^* \equiv (\Delta_1^*, \Delta_2^*)$ be the true parameter value

Suppose $S = S_I$, then a BNE under S_I induces the following CCP:

$$P_{10} \equiv P(y = (1,0)) = \frac{1 + \Delta_2^*}{(2 + \Delta_1^*)(2 + \Delta_2^*)}$$

- e.g., if $\Delta_1^* = \Delta_2^* = 0.5$, then $P_{10} = 0.24 < LB_\Delta$ for all Δ

Hence, we can detect the violation of H_0 by comparing the CCP and BCE prediction (i.e., LB_Δ)

Support Function

The support function

$$h(b, A) \equiv \sup_{q \in A} b'q, \;\; b \in \mathbb{B}_x$$

is a continuous function on the "unit ball":

$$\mathbb{B}_x \equiv \{b \in \mathbb{R}^{|Y|} : \|b\|_{W_x} \leq 1\}, \;\; \|b\|_{W_x} = (b' W_x b)^{1/2}$$

where $W_x \equiv \mathrm{AsyVar}(\hat{P}_{n,x})$

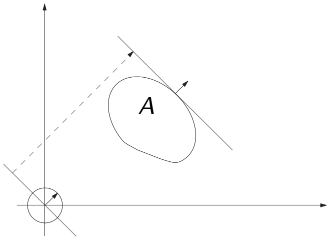

Computation

A key object is

$$V_{n,x}(\theta) \equiv \sup_{b \in \mathbb{B}_{n,x}} \sqrt{n}\{b'\hat{P}_{n,x} - h(b, Q^{BCE}_{\theta,S^r}(x))\}$$

$$= \sup_{b \in \mathbb{B}_{n,x}} \inf_{q \in Q^{BCE}_{\theta,S^r}(x)} \sqrt{n}\left[b'\hat{P}_{n,x} - b'q\right] \quad (P0)$$

Problem (P0) can be recast as a convex quadratic program:

$$V_{n,x}(\theta) = \max_{lb \leq w \leq ub} -\gamma' w$$

$$s.t. w'\Gamma_1 w \leq 1$$

$$\Gamma_2 w = 0_{|Y|}$$

$$\Gamma_3 w \leq 0_{d_\nu}$$

- $w \equiv (b', \lambda'_{eq}, \lambda'_{ineq})'$ stacks $b \in \mathbb{R}^{|Y|}$ and Lagrange multipliers associated with the constraints

Return

Confidence Set for x Satisfying Null

$$H_{0,x} : S(x) \succeq S_{Priv,Hub}(x)$$

Let X_0 be the set of x's for which $H_{0,x}$ is true

Define

$$T_x(\theta) \equiv \sup_{b \in \mathbb{B}_{n,x}} \{b' P_{y|x} - h(b, Q^{BCE}_{\theta,S^r}(x))\}$$

Define the bootstrap p-value as

$$p_n(\theta, x) \equiv P^*(T^*_{n,x}(\theta) > T_{n,x}(\theta)|y^n, x^n)$$

- $T_{n,x}$ and $T^*_{n,x}$ are sample and bootstrap analogs of T_x

Then define a confidence set for X_0 as:

$$CS_n \equiv \{x : p_n(x) > \alpha_x\}$$

where α_x is chosen to control for FWER