Counterfactual Sensitivity and Robustness[*]

Timothy Christensen[†]   Benjamin Connault[‡]

May 25, 2022

Abstract

We propose a framework for analyzing the sensitivity of counterfactuals to parametric assumptions about the distribution of latent variables in structural models. In particular, we derive bounds on counterfactuals as the distribution of latent variables spans nonparametric neighborhoods of a given parametric specification while other "structural" features of the model are maintained. Our approach recasts the infinite-dimensional problem of optimizing the counterfactual with respect to the distribution of latent variables (subject to model constraints) as a finite-dimensional convex program. We also develop an MPEC version of our method to further simplify computation in models with endogenous parameters (e.g., value functions) defined by equilibrium constraints. We propose plug-in estimators of the bounds and two methods for inference. We also show that our bounds converge to the sharp nonparametric bounds on counterfactuals as the neighborhood size becomes large. To illustrate the broad applicability of our procedure, we present empirical applications to matching models with transferable utility and dynamic discrete choice models.

Keywords: Robustness, ambiguity, model uncertainty, misspecification, global sensitivity analysis.

JEL codes: C14, C18, C54, D81

1 Introduction

Researchers frequently make parametric assumptions about the distribution of latent variables in structural models. These assumptions are typically made for computational convenience[1] or because simulation-based methods are used for estimation. In many models, such as those we consider in this paper, the distribution of latent variables is not nonparametrically identified. This raises the possibility that model parameters and the outcomes of policy experiments, or *counterfactuals*, may be only partially identified when parametric assumptions are relaxed. That is, different distributions may fit the data equally well in-sample, but may yield different values of the counterfactual. It is therefore natural to question whether counterfactuals are sensitive or robust to researchers' parametric assumptions, especially when evaluating the credibility of structural modeling exercises.

This paper proposes a framework for analyzing the sensitivity of counterfactuals to parametric assumptions about the distribution of latent variables in a class of structural models. In particular, we derive bounds on counterfactuals as the distribution of latent variables spans nonparametric neighborhoods of a given parametric specification while other "structural" features of the model are maintained. This approach is in the spirit of global sensitivity analysis advocated by Leamer (1985) (see also Tamer (2015)). Global sensitivity analyses are important in this context: many structural models are nonlinear so policy interventions can have different effects at different points in the parameter space. But a major difficulty with implementing global sensitivity analyses is tractability. A more tractable alternative are local sensitivity analyses, which are based on small perturbations around a chosen specification. Because local approaches rely on linearization, they may fail to correctly characterize the range of counterfactuals predicted by a nonlinear model when the distribution differs nontrivially from the researcher's chosen parametric specification.

Our main insight is to borrow from the robustness literature in economics pioneered by Hansen and Sargent (2001, 2008) to simplify computation using convex programming.[2] Following this literature, we define neighborhoods around the researcher's para-

---

[1] Examples include the conventional Gumbel (or type-I extreme value) assumption in discrete choice models following McFadden (1974), dynamic discrete choice models following Rust (1987), and matching models with transferable utility following Dagsvik (2000) and Choo and Siow (2006). Models of static or dynamic discrete games often impose parametric assumptions about the distribution of payoff shocks—see, e.g., Berry (1992), Aguirregabiria and Mira (2007), Bajari, Benkard, and Levin (2007), and Ciliberto and Tamer (2009).

[2] Our approach is also related to the field of *distributionally robust optimization* in operations re-

metric specification using statistical divergence (e.g., Kullback–Leibler divergence), with the option to add certain shape restrictions as appropriate. For tractability, we restrict our attention to models that may be written as a finite number of moment (in)equalities, where the expectation is with respect to the distribution of latent variables. While restrictive, this class accommodates many important models of static and dynamic discrete choice, discrete games, and matching.

To describe our procedure, consider the problem of minimizing or maximizing the counterfactual at a fixed value of structural parameters by varying the distribution of latent variables over a neighborhood, subject to the model's (in)equality restrictions. We use duality to recast this infinite-dimensional optimization problem as a finite-dimensional convex program. The value of this *inner* program is treated as a criterion function, which is optimized in an *outer* optimization with respect to structural parameters. Importantly, the dimension of the inner problem is independent of the neighborhood size, making our procedure tractable over both small and large neighborhoods. To further simplify computation, we develop an MPEC version of our procedure for models featuring endogenous parameters (e.g., value functions) defined by equilibrium constraints. We show that this implementation can produce significant computational gains for dynamic discrete choice models in particular.

Our approach is conceptually different from nonparametric partial identification analyses which derive bounds on counterfactuals under minimal distributional assumptions. But as we show, bounds computed using our procedure converge to the (sharp) nonparametric bounds in the limit as the neighborhood size becomes large. Aside from sensitivity analyses, our methods may therefore be used to approximate nonparametric bounds by taking the neighborhood size to be large but finite.

For estimation and inference, we propose simple plug-in estimators of the bounds and establish their consistency. We also propose and theoretically justify two methods for inference: a computationally simple but conservative projection procedure and a relatively more efficient bootstrap procedure.

We illustrate our procedures with two empirical applications. The first revisits the "marital college premium" estimates reported in Chiappori, Salanié, and Weiss (2017), which relied on an i.i.d. Gumbel (type-I extreme value) assumption for the distribution of individuals' idiosyncratic marital preferences (see also Choo and Siow (2006)). The second empirical application performs a counterfactual welfare analysis in the canonical dynamic discrete choice model of Rust (1987).

---

search. See, e.g., Shapiro (2017), Duchi and Namkoong (2021), and references therein.

Related literature. Our approach has connections with global prior sensitivity in Bayesian analysis (Chamberlain and Leamer, 1976; Leamer, 1982; Berger, 1984), most notably Giacomini, Kitagawa, and Uhlig (2016) and Ho (2018) who consider sets of priors constrained by Kullback–Leibler divergence relative to a default prior.

Motivated by questions of sensitivity, Chen, Tamer, and Torgovitsky (2011) study inference in semiparametric likelihood models using sieve approximations for the infinite-dimensional nuisance parameter (the distribution of latent variables in our setting). For the class of moment-based models we consider, our approach instead eliminates the infinite-dimensional nuisance parameter via a convex program of fixed dimension.

Several other works have used convex duality to characterize identified sets in models with latent variables. Most closely related are Ekeland, Galichon, and Henry (2010) and Schennach (2014).[3] The problem we study is different, both because of its focus on counterfactuals, rather than structural parameters, and because the optimization is performed over a neighborhood, rather than over all distributions. As a consequence, our estimation and inference methods are also quite different.

Torgovitsky (2019b) uses linear programming to characterize sharp identified sets in latent variable models defined by quantile restrictions. Within this class, his approach is more computationally convenient than ours for characterizing identified sets. Several important moments or counterfactuals cannot be expressed as quantile restrictions, such as social surplus in discrete choice models and Bellman equations in dynamic discrete choice models. Our approach is compatible with these moments and counterfactuals, thereby allowing the user to characterize identified sets in broader classes of model as well as to perform sensitivity analyses.

There is also a literature deriving nonparametric bounds in specific latent variable models. Examples include Manski (2007, 2014), Allen and Rehbeck (2019), Tebaldi, Torgovitsky, and Yang (2019), Lafférs (2019), Torgovitsky (2019a), and Gualdani and Sinha (2020). Most closely related is Norets and Tang (2014), who construct identified sets of counterfactual conditional choice probabilities (CCPs) in dynamic binary choice models. Their approach is specific to counterfactual CCPs and to dynamic binary choice models. Our approach allows for a wider range of counterfactual (e.g., welfare), shape restrictions, and multinomial choice, in addition to performing sensitivity analyses.[4]

---

[3]Works using other notions of "duality" to construct identified sets include Beresteanu, Molchanov, and Molinari (2011), Galichon and Henry (2011), Chesher and Rosen (2017), and Li (2018).

[4]Kalouptsidi, Scott, and Souza-Rodrigues (2021) and Kalouptsidi, Kitamura, Lima, and Souza-Rodrigues (2020) consider the converse problem, in which flow payoffs are nonparametric (as they can be in our setting) but the distribution of latent payoff shocks is known.

Finally, our work is complementary to the recent literature on local sensitivity—see, e.g., Kitamura, Otsu, and Evdokimov (2013), Andrews, Gentzkow, and Shapiro (2017, 2020), Armstrong and Kolesár (2021), Bonhomme and Weidner (2021), and Mukhin (2018). Much of this literature is concerned with local misspecification of moment conditions, which is different from the setting we consider.

Outline. Section 2 introduces our procedure, estimators of the bounds, and shows our approach recovers nonparametric bounds as the neighborhood size becomes large. Section 3 discusses practical aspects and implementation details. Section 4 gives guidance for interpreting the neighborhood size. Empirical applications are presented in Section 5. Section 6 discusses estimation and inference. The online appendix presents extensions of our methodology, connections with local sensitivity analyses, additional empirical results, and proofs of our main results. A secondary appendix of our working paper version Christensen and Connault (2022) presents background material on Orlicz classes and supplemental proofs.

2 Procedure

We begin in Section 2.1 by describing the class of models to which our procedure may be applied. Section 2.2 describes our approach, Section 2.3 shows how duality is used to simplify the bounds, and Section 2.4 introduces our estimators of the bounds. Section 2.5 shows our bounds converge to the sharp nonparametric bounds as the neighborhood size becomes large.

2.1 Setup

We consider a class of models that link a structural parameter $\theta \in \Theta \subset \mathbb{R}^{d_\theta}$, a vector of targeted moments $P_0 \in \mathcal{P} \subseteq \mathbb{R}^{d_P}$, and possibly an auxiliary parameter $\gamma_0 \in \Gamma$ (a metric space) via the moment restrictions

$$\mathbb{E}^F[g_1(U, \theta, \gamma_0)] \leq P_{10}, \tag{1a}$$

$$\mathbb{E}^F[g_2(U, \theta, \gamma_0)] = P_{20}, \tag{1b}$$

$$\mathbb{E}^F[g_3(U, \theta, \gamma_0)] \leq 0, \tag{1c}$$

$$\mathbb{E}^F[g_4(U, \theta, \gamma_0)] = 0, \tag{1d}$$

where $g_1, \ldots, g_4$ are vectors of moment functions, $P_0 = (P_{10}, P_{20})$ is partitioned conformably, and $\mathbb{E}^F$ denotes expectation with respect to a vector of latent variables $U \sim F$. We assume that the researcher has consistent estimators $(\hat{P}, \hat{\gamma})$ of (P_0, γ_0). We also assume that the researcher is interested in a (scalar) counterfactual of the form

$$\kappa = \mathbb{E}^F[k(U, \theta, \gamma_0)]. \tag{2}$$

This setup accommodates counterfactuals that do not depend explicitly on U, in which case (2) reduces to $\kappa = k(\theta, \gamma_0)$. Note that κ will still depend on the distribution of U through θ, whose values are disciplined by the moment conditions (1).

Several models and counterfactuals of interest fall into this framework. We review three examples before proceeding.

Example 2.1 (Discrete choice and consumer welfare) Suppose an individual derives utility $h_j(X, \theta) + U_j$ from choice $j \in \mathcal{J}_0 := \{0, 1, \ldots, J\}$, where $X \in \mathcal{X}$ are observed covariates and $U = (U_j)_{j \in \mathcal{J}_0}$ is latent (to the econometrician). We assume, as typical, that U is drawn independently across individuals from a continuous distribution F. The probability that an individual with characteristics x chooses j is

$$p(j|x) = \mathbb{P}_F \left(h_j(x, \theta) + U_j = \max_{j' \in \mathcal{J}_0} \left(h_{j'}(x, \theta) + U_{j'} \right) \right), \tag{3}$$

where $\mathbb{P}_F$ denotes probabilities when $U \sim F$. In empirical work, θ is typically estimated using a criterion that fits the model-implied choice probabilities (3) to probabilities observed in the data. Welfare analyses are often based on the social surplus (McFadden, 1978)

$$W(x) = \mathbb{E}^F \left[\max_{j \in \mathcal{J}_0} \left(h_j(x, \theta) + U_j \right) \right],$$

which is the average utility consumers with characteristics x derive from the choice problem. A related welfare measure is the change in surplus $\Delta W(x_a, x_b) = W(x_a) - W(x_b)$ associated with a shift from x_b to x_a. In practice, it is common to assume the U_j are i.i.d. Gumbel (type-I extreme value), as this yields closed-form expressions for choice probabilities and the welfare measures $W(x)$ and $\Delta W(x_a, x_b)$.

Our approach may be used to perform a sensitivity analysis of $W(x)$ and $\Delta W(x_a, x_b)$ to parametric assumptions about F when $\mathcal{X}$ is finite. A leading example is matching models with finitely many agent types—see Section 5.1 and references therein. Understanding the sensitivity of $W(x)$ and $\Delta W(x_a, x_b)$ to F is important in this case because

$W(x)$ and $\Delta W(x_a, x_b)$ are not nonparametrically identified.[5]

In our notation, g_2 collects indicator functions representing the choice probabilities (3) across covariates $x \in \mathcal{X}$ and choices $j \in \mathcal{J} := \{1, \ldots, J\}$ ($j = 0$ is redundant):

$$g_2(U, \theta) = (\mathbb{1}\{h_j(x, \theta) + U_j = \max_{j' \in \mathcal{J}_0}(h_{j'}(x, \theta) + U_{j'})\})_{(j,x) \in \mathcal{J} \times \mathcal{X}}$$

and $P_{20} = (\Pr(j|x))_{(j,x) \in \mathcal{J} \times \mathcal{X}}$ is the vector of true choice probabilities. There are no g_1, g_3, g_4, or γ in this model. Finally, $k(U, \theta) = \max_{j \in \mathcal{J}_0}(h_j(x, \theta) + U_j)$ for $W(x)$ and $k(U, \theta) = \max_{j \in \mathcal{J}_0}(h_j(x_a, \theta) + U_j) - \max_{j \in \mathcal{J}_0}(h_j(x_b, \theta) + U_j)$ for $\Delta W(x_a, x_b)$. $\qquad\square$

Example 2.2 (Discrete games) Following Bresnahan and Reiss (1990, 1991), Berry (1992), and Tamer (2003), consider the complete-information game in Table 1.

		Firm 2	
		0	1
Firm 1	0	$(0,0)$	$(0, \beta_2'x + U_2)$
	1	$(\beta_1'x + U_1, 0)$	$(\beta_1'x - \Delta_1 + U_1, \beta_2'x - \Delta_2 + U_2)$

Table 1: Payoff matrix for (Firm 1, Firm 2) when $X = x$.

Here $U = (U_1, U_2)$ is the latent (to the econometrician) component of firms' profits, which is independent of covariates X. Suppose that the solution concept is restricted to equilibria in pure strategies. The econometrician may estimate the probabilities of the potential market structures $(0,0)$, $(0,1)$, $(1,0)$, $(1,1)$ (conditional on X) from data on a large number of markets. As the model is incomplete—there are values of U for which there are multiple equilibria—moment inequality methods are typically used in empirical work to avoid restricting the equilibrium selection mechanism. However, strong parametric assumptions are often made about the distribution of U (typically bivariate Normal) to derive the model-implied probabilities for different market structures; see, e.g., Berry (1992), Ciliberto and Tamer (2009), Beresteanu et al. (2011), and Kline and Tamer (2016). It therefore seems natural to also question the sensitivity of counterfactuals to parametric assumptions for U.

This model falls into our setup when the regressors X have finite support $\mathcal{X}$.[6] In

---

[5]See, e.g., Berry and Haile (2010, 2014) and Allen and Rehbeck (2019) for nonparametric identification of utilities and welfare measures in discrete choice models when characteristics have continuous support.

[6]Continuous regressors are often discretized in empirical applications; see, e.g., Ciliberto and Tamer (2009), Grieco (2014), Kline and Tamer (2016), and Chen, Christensen, and Tamer (2018).

our notation, g_1 collects the moment inequalities that bound the probabilities of $(0, 1)$ and $(1, 0)$ across $x \in \mathcal{X}$, with P_{10} denoting the corresponding true probabilities. The inequalities are typically expressed as upper bounds on the probabilities of $(0, 1)$ and $(1, 0)$; we flip the sign to be compatible with (1a):

$$g_1(U, \theta) = \begin{bmatrix} (-1\!\!1\{U_1 \geq -\beta_1' x; U_2 \leq \Delta_2 - \beta_2' x\})_{x \in \mathcal{X}} \\ (-1\!\!1\{U_1 \leq \Delta_1 - \beta_1' x; U_2 \geq -\beta_2' x\})_{x \in \mathcal{X}} \end{bmatrix}, \quad P_{10} = \begin{bmatrix} (-\Pr((1, 0)|X = x))_{x \in \mathcal{X}} \\ (-\Pr((0, 1)|X = x))_{x \in \mathcal{X}} \end{bmatrix},$$

where $\theta = (\Delta_1, \Delta_2, \beta_1, \beta_2)$. Similarly, g_2 and P_{20} collect the moment conditions and probabilities for outcomes $(0, 0)$ and $(1, 1)$, which are always realized as the result of unique equilibria:

$$g_2(U, \theta) = \begin{bmatrix} (1\!\!1\{U_1 \leq -\beta_1' x;\ U_2 \leq -\beta_2' x\})_{x \in \mathcal{X}} \\ (1\!\!1\{U_1 \geq \Delta_1 - \beta_1' x;\ U_2 \geq \Delta_2 - \beta_2' x\})_{x \in \mathcal{X}} \end{bmatrix}, \quad P_{20} = \begin{bmatrix} (\Pr((0, 0)|X = x))_{x \in \mathcal{X}} \\ (\Pr((1, 1)|X = x))_{x \in \mathcal{X}} \end{bmatrix}.$$

There is no g_3, g_4, or γ in this model. Ciliberto and Tamer (2009) compute upper bounds on the probability of entrants under a counterfactual payoff shift, say $\tau(\theta)$. The function $k(U, \theta) = 1\!\!1\{U_1 \geq \tau(\theta) - \beta_1' x\}$ corresponds to the upper bound on the probability of firm 1 entering when $X = x$ under this counterfactual. $\square$

Example 2.3 (Dynamic discrete choice) Consider a canonical dynamic discrete choice (DDC) model following Rust (1987). The decision maker solves

$$V(s) = \mathbb{E}^F \left[\max_{d \in \mathcal{D}_0} \left(\pi_{d,s}(\theta_\pi) + U_d + \beta E[V(s')|d, s] \right) \right], \tag{4}$$

where $s \in \mathcal{S}$ is a Markov state variable, $\mathcal{D}_0 = \{0, 1, \ldots, D\}$ is the set of actions, $\pi_{d,s}$ is the flow payoff for action d in state s which is parameterized by θ_π, U_d is a latent payoff shock, $\beta \in (0, 1)$ is a discount parameter, and $E[\,\cdot\,|d, s]$ denotes expectation with respect to the future state s'. The distribution F of $U = (U_d)_{d \in \mathcal{D}_0}$ is typically assumed to be continuous and independent of s. The CCP of action d in state s is

$$p(d|s) = \mathbb{P}_F \left(\pi_{d,s}(\theta_\pi) + U_d + \beta E[V(s')|d, s] = \max_{d' \in \mathcal{D}_0} \left(\pi_{d',s}(\theta_\pi) + U_{d'} + \beta E[V(s')|d', s] \right) \right), \tag{5}$$

where $\mathbb{P}_F$ denotes probabilities when $U \sim F$.

It is standard to assume the U_d are i.i.d. Gumbel, as this yields closed-form expressions for the expectation in (4) and multinomial-logit expressions for the CCPs (5).

Parameters θ_π or (θ_π, β) are typically estimated using a criterion function that fits the model-implied CCPs (5) to probabilities observed in the data. Counterfactuals are then computed by solving (4) under alternative laws of motion, flow payoffs, or other interventions.

When $\mathcal{S}$ is finite, model parameters, counterfactual CCPs, and counterfactual welfare measures are typically not identified without parametric restrictions on F. Our procedure may be used perform a sensitivity analysis of counterfactuals to parametric assumptions on F as follows. Let $\theta = (\theta_\pi, v, \tilde{v})$ or $\theta = (\theta_\pi, \beta, v, \tilde{v})$, where $v = (V(s))_{s \in \mathcal{S}}$ and $\tilde{v} = (\tilde{V}(s))_{s \in \mathcal{S}}$ collect the baseline and counterfactual value functions across $s \in \mathcal{S}$. Also let $\gamma = (M_d)_{d \in \mathcal{D}_0}$ collect the transition matrices for s, g_2 collect indicator functions for the CCPs (5) across states $s \in \mathcal{S}$ and choices $d \in \mathcal{D} := \{1, \ldots, D\}$ ($d = 0$ is redundant):

$$g_2(U, \theta, \gamma) = \left(\mathbb{1}\left\{ \pi_{d,s}(\theta_\pi) + U_d + \beta M_{d,s} v = \max_{d' \in \mathcal{D}_0} \left(\pi_{d',s}(\theta_\pi) + U_{d'} + \beta M_{d',s} v \right) \right\} \right)_{(d,s) \in \mathcal{D} \times \mathcal{S}}$$

with $M_{d,s}$ denoting the sth row of M_d, and $P_{20} = (\Pr(d|s))_{(d,s) \in \mathcal{D} \times \mathcal{S}}$ collect the corresponding true CCPs. Finally, g_4 collects moment functions representing (4) in the baseline model and under the counterfactual:

$$g_4(U, \theta, \gamma) = \begin{bmatrix} (\max_{d \in \mathcal{D}_0}\{\pi_{d,s}(\theta_\pi) + U_d + \beta M_{d,s} v\} - v_s)_{s \in \mathcal{S}} \\ (\max_{d \in \mathcal{D}_0}\{\tilde{\pi}_{d,s}(\theta_\pi) + U_d + \tilde{\beta} \tilde{M}_{d,s} \tilde{v}\} - \tilde{v}_s)_{s \in \mathcal{S}} \end{bmatrix}, \quad (6)$$

where $v_s = V(s)$, $\tilde{v}_s = \tilde{V}(s)$, and $\tilde{\pi}$, $\tilde{\beta}$, $\tilde{M}_d$ denote counterfactual flow payoffs, discount factor, and law of motion.[7] We recommend including the location normalizations $\mathbb{E}^F[U_d] = 0$ for $d \in \mathcal{D}_0$ in g_4 for interpretability. We also recommend including scale normalizations in g_4 so that $\mathbb{E}^F[\max_{d \in \mathcal{D}_0} U_d]$ is finite. For instance, in Section 5.2 we normalize $\mathbb{E}^F[U_d^2]$ for all $d \in \mathcal{D}_0$.

Counterfactual CCPs can be computed using

$$k(U, \theta, \gamma) = \mathbb{1}\left\{ \tilde{\pi}_{d,s}(\theta_\pi) + U_d + \tilde{\beta} \tilde{M}_{d,s} \tilde{v} = \max_{d' \in \mathcal{D}_0} \left(\tilde{\pi}_{d',s}(\theta_\pi) + U_{d'} + \tilde{\beta} \tilde{M}_{d',s} \tilde{v} \right) \right\}.$$

Change in average welfare corresponds to $k(\theta, \gamma) = w'(\tilde{v} - v)$ for a weight vector w. $\square$

---

[7]If $\mathbb{E}^F[\max_{d \in \mathcal{D}_0} U_d]$ is finite, then $v \mapsto (\mathbb{E}^F[\max_{d \in \mathcal{D}_0}\{\pi_{d,s}(\theta_\pi) + U_d + \beta M_{d,s} v\}])_{s \in \mathcal{S}}$ is a ℓ^∞-contraction of modulus β on $\mathbb{R}^{|\mathcal{S}|}$. Hence, there is a unique $(v, \tilde{v})$ solving $\mathbb{E}^F[g_4(U, \theta, \gamma)] = 0$ at any fixed $(\theta_\pi, \beta, \tilde{\beta}, F)$. The solution $(v, \tilde{v})$ must collect the solutions to (4) in the baseline model and counterfactual across states: $v = (V(s))_{s \in \mathcal{S}}$ and $\tilde{v} = (\tilde{V}(s))_{s \in \mathcal{S}}$. It follows that F satisfies $\mathbb{E}^F[g_4(U, \theta, \gamma)] = 0$ at $\theta = (\theta_\pi, \beta, v, \tilde{v})$ if and only if $(v, \tilde{v})$ corresponds to the value functions V and $\tilde{V}$ under F.

Remark 2.1 *We allow for conditional moments models with* $\mathbb{E}[g_1(U, X, \theta, \gamma)|X = x] \leq P_{10}(x)$ *(and similarly for (1b)-(1d)) if U is independent of X and X takes values in a finite set $\mathcal{X}$. Moment functions are then stacked across $x \in \mathcal{X}$ to form g_1, g_2, g_3, and g_4 (see Examples 2.1-2.3). Appendix A discusses extensions to conditional moment models where the distribution of U may vary with the value of (discrete) covariates, and to non-separable models with discrete covariates. Models with continuous covariates fall outside the scope of our procedure.*

Remark 2.2 *Our setup relies on the counterfactual being expressible as (2). If k is vector-valued, our procedure can be applied to compute the support function[8] of the identified set of counterfactuals: set $k^\tau(U, \theta, \gamma) = \tau' k(U, \theta, \gamma)$ for a conformable unit vector τ and replace (2) with $\kappa^\tau = \mathbb{E}^F[k^\tau(U, \theta, \gamma_0)]$. Our setup excludes counterfactuals that are infinite-dimensional, such as the distribution of the number of firms in a market.*

Remark 2.3 *The distribution F is not nonparametrically identified in any of the above examples or, more generally, in the class of models (1) when the support of U contains many more points than there are moment conditions (e.g., when U is continuously distributed).*

In common practice, a seemingly reasonable or computationally convenient distribution, say F_*, is assumed by the researcher and maintained throughout the analysis (e.g., bivariate Normal in Example 2.2 and i.i.d. Gumbel in Examples 2.1 and 2.3). Given F_* and estimates $\hat{P} = (\hat{P}_1, \hat{P}_2)$ of P_0 and $\hat{\gamma}$ of γ_0, the researcher computes an estimate $\hat{\theta}$ of θ using a criterion function based on the moment conditions

$$
\begin{aligned}
\mathbb{E}^{F_*}[g_1(U, \theta, \hat{\gamma})] \leq \hat{P}_1, \quad &\mathbb{E}^{F_*}[g_2(U, \theta, \hat{\gamma})] = \hat{P}_2, \\
\mathbb{E}^{F_*}[g_3(U, \theta, \hat{\gamma})] \leq 0, \quad &\mathbb{E}^{F_*}[g_4(U, \theta, \hat{\gamma})] = 0.
\end{aligned}
\tag{7}
$$

Finally, the researcher estimates the counterfactual using $\hat{\kappa} = \mathbb{E}^{F_*}[k(U, \hat{\theta}, \hat{\gamma})]$. If k does not depend on U, then the estimated counterfactual is simply $\hat{\kappa} = k(\hat{\theta}, \hat{\gamma})$. In this case $\hat{\kappa}$ will still depend implicitly on F_* through $\hat{\theta}$.[9]

The researcher's chosen specification F_* is used both for estimation of θ and again when computing the counterfactual. A natural question is: to what extent does the counterfactual depend on the choice of distribution? The main contribution of this paper is to provide a tractable econometric framework for answering this question.

---

[8]A closed convex set is determined by its support function—see Rockafellar (1970, Section 13).

[9]While this discussion has assumed point identification of θ and κ for sake of exposition, our methods allow structural parameters and counterfactuals to be partially identified.

2.2 Our Approach

As a sensitivity analysis, we shall relax the researcher's parametric assumption and allow F to vary over nonparametric neighborhoods $\mathcal{N}_\delta$ of F_*, where δ is a measure of neighborhood "size". When we do so, there may be multiple pairs $(\theta, F) \in \Theta \times \mathcal{N}_\delta$ that satisfy (1) but which yield different values of the counterfactual. Our objects of interest are the smallest and largest values of the counterfactual over all such (θ, F) pairs:

$$\underline{\kappa}_\delta = \inf_{\theta \in \Theta, F \in \mathcal{N}_\delta} \mathbb{E}^F[k(U, \theta, \gamma_0)] \quad \text{subject to (1),} \tag{8}$$

$$\overline{\kappa}_\delta = \sup_{\theta \in \Theta, F \in \mathcal{N}_\delta} \mathbb{E}^F[k(U, \theta, \gamma_0)] \quad \text{subject to (1).} \tag{9}$$

By focusing on $\underline{\kappa}_\delta$ and $\overline{\kappa}_\delta$, our approach naturally accommodates models with partially-identified structural parameters and counterfactuals. Our approach also sidesteps having to compute the identified set of structural parameters.

The optimization problems (8) and (9) are made tractable by a convenient choice of $\mathcal{N}_\delta$. Following Hansen and Sargent (2001) and Maccheroni, Marinacci, and Rustichini (2006), we consider neighborhoods constrained by ϕ-divergence (Csiszár, 1975):

$$\mathcal{N}_\delta = \{F \in \mathcal{F} : D_\phi(F \| F_*) \leq \delta\},$$

$$D_\phi(F \| F_*) = \begin{cases} \int \phi\left(\frac{\mathrm{d}F}{\mathrm{d}F_*}\right) \mathrm{d}F_* & \text{if } F \ll F_*, \\ +\infty & \text{otherwise,} \end{cases} \tag{10}$$

where $\mathcal{F}$ denotes all probability measures on the support[10] $\mathcal{U}$ of U and $F \ll F_*$ denotes absolute continuity of F with respect to F_*. The convex function $\phi : [0, \infty) \to \mathbb{R}_+ \cup \{+\infty\}$ penalizes deviations of F from F_*. For example, $\phi(x) = x \log x - x + 1$ corresponds to Kullback–Leibler (KL) divergence, $\phi(x) = \frac{1}{2}(x - 1)^2$ corresponds to Pearson χ^2 divergence, and

$$\phi(x) = \frac{x^p - 1 - p(x - 1)}{p(p - 1)}, \quad (p > 1),$$

corresponds to L^p divergence. If F_* has positive (Lebesgue) density, then the absolute continuity condition merely rules out F with mass points.

Remark 2.4 *Normalizations and other shape restrictions may be added by augmenting the moment functions $g_1, \ldots, g_4$. Examples include: (i) location normalizations, e.g.*

---

[10]That is, $\mathcal{U}$ is the set of all values that U could conceivably take according to the model, which is possibly larger that the support of the measure F_*.

$\mathbb{E}^F[U] = 0$ or $\mathbb{E}^F[\mathbb{1}\{U_i \leq 0\} - 0.5] = 0$ *for each element U_i of U; (ii) scale normalizations, e.g. $\mathbb{E}^F[U_i^2] = 1$; (iii) covariance normalizations, e.g. $\mathbb{E}^F[UU'] = I$; and (iv) smoothness restrictions, e.g. $\mathbb{E}^F[\mathbb{1}\{U_i \leq a_{k+1}\} - \mathbb{1}\{U_i \leq a_k\}] \leq C$ for $a_1 < \ldots < a_K$ and a positive constant C.*

Remark 2.5 *Appendix A.1 shows that shape restrictions including symmetry, exchangeability, and, more generally, invariance under a finite group of transforms, are also easy to impose.*

2.3 Dual Formulation

We use convex duality to simplify computation of $\underline{\kappa}_\delta$ and $\overline{\kappa}_\delta$. We start by noting $\underline{\kappa}_\delta$ and $\overline{\kappa}_\delta$ may be written as the solution to two profiled optimization problems:

$$\underline{\kappa}_\delta = \inf_{\theta \in \Theta} \underline{K}_\delta(\theta; \gamma_0, P_0), \qquad\qquad \overline{\kappa}_\delta = \sup_{\theta \in \Theta} \overline{K}_\delta(\theta; \gamma_0, P_0),$$

where the criterion functions $\underline{K}_\delta(\theta; \gamma_0, P_0)$ and $\overline{K}_\delta(\theta; \gamma_0, P_0)$ are, respectively, the infimum and supremum of $\mathbb{E}^F[k(U, \theta, \gamma_0)]$ with respect to $F \in \mathcal{N}_\delta$ subject to the moment conditions (1). In what follows, it is helpful to define the criterion functions at a generic (γ, P). To do so, we say that the moment conditions (1) hold "at (θ, γ, P)" if they hold when γ_0 is replaced by γ and P_0 is replaced by P. Then

$$\underline{K}_\delta(\theta; \gamma, P) = \inf_{F \in \mathcal{N}_\delta} \mathbb{E}^F[k(U, \theta, \gamma)] \quad \text{subject to (1) holding at } (\theta, \gamma, P), \qquad (11)$$

$$\overline{K}_\delta(\theta; \gamma, P) = \sup_{F \in \mathcal{N}_\delta} \mathbb{E}^F[k(U, \theta, \gamma)] \quad \text{subject to (1) holding at } (\theta, \gamma, P), \qquad (12)$$

with the understanding that $\underline{K}_\delta(\theta; \gamma, P) = +\infty$ and $\overline{K}_\delta(\theta; \gamma, P) = -\infty$ if there does not exist a distribution in $\mathcal{N}_\delta$ for which the moment conditions (1) hold at (θ, γ, P).

We first impose some mild regularity conditions on F_*, ϕ, and the moment functions to justify the dual formulation. Similar conditions are used in generalized empirical likelihood estimation (see, e.g., Komunjer and Ragusa (2016)). Let Φ_0 denote the set of all $\phi : [0, \infty) \to \mathbb{R} \cup \{+\infty\}$ such that ϕ is continuously differentiable on $(0, +\infty)$ and strictly convex, with $\phi(1) = \phi'(1) = 0$, $\phi(0) < +\infty$, $\lim_{x \downarrow 0} \phi'(x) < 0$, $\lim_{x \to +\infty} \phi(x)/x = +\infty$, $\lim_{x \to +\infty} \phi'(x) > 0$, and $\lim_{x \to +\infty} x\phi'(x)/\phi(x) < +\infty$. The functions inducing KL, χ^2, and L^p divergence all belong to Φ_0.

Let $\phi^\star(x) = \sup_{t \geq 0 : \phi(t) < +\infty}(tx - \phi(t))$ denote the convex conjugate of $\phi \in \Phi_0$ and let $\psi(x) = \phi^\star(x) - x$. Define $\mathcal{E} = \{f : \mathcal{U} \to \mathbb{R} \text{ for which } \mathbb{E}^{F_*}[\psi(c|f(U)|)] < \infty \text{ for all } c > 0\}$.

The class $\mathcal{E}$ is an Orlicz class of functions (see Appendix F of Christensen and Connault (2022) for details). For example,

$$\mathcal{E} = \{f : \mathcal{U} \to \mathbb{R} : \mathbb{E}^{F_*}[e^{c|f(U)|}] < \infty \text{ for all } c > 0\} \quad \text{for KL divergence,}$$
$$\mathcal{E} = \{f : \mathcal{U} \to \mathbb{R} : \mathbb{E}^{F_*}[f(U)^2] < \infty\} \quad \text{for } \chi^2 \text{ divergence, and}$$
$$\mathcal{E} = \{f : \mathcal{U} \to \mathbb{R} : \mathbb{E}^{F_*}[|f(U)|^q] < \infty\} \quad \text{for } L^p \text{ divergence } (p^{-1} + q^{-1} = 1).$$

Let $g = (g_1, g_2, g_3, g_4)$ denote the vector formed by stacking each of the moment functions from (1a)–(1d). Our key regularity condition is the following:

Assumption Φ *(i) $\phi \in \Phi_0$.*

(ii) $k(\,\cdot\,, \theta, \gamma)$ and each entry of $g(\,\cdot\,, \theta, \gamma)$ belong to $\mathcal{E}$ for each $\theta \in \Theta$ and $\gamma \in \Gamma$.

For KL divergence, the class $\mathcal{E}$ contains bounded functions (e.g., indicator functions) and functions that are additively separable in U provided F_* has tails that decay faster than exponentially (e.g., Gaussian but not Gumbel). Assumption Φ therefore fails for KL divergence in Examples 2.1 and 2.3, but holds for χ^2 or L^p divergence as these only require finite second or qth moments, respectively.

Let $d = \sum_{i=1}^{4} d_i$ where d_i is the dimension of g_i, let $\Lambda = \mathbb{R}_+^{d_1} \times \mathbb{R}^{d_2} \times \mathbb{R}_+^{d_3} \times \mathbb{R}^{d_4}$, and let λ_{12} denote the first $d_1 + d_2$ elements of λ. A derivation of the following criterion functions is presented in Appendix G.2 of Christensen and Connault (2022).

Proposition 2.1 *Suppose that Assumption Φ holds. Then the criterion functions (11) and (12) may be restated as*

$$\underline{K}_\delta(\theta; \gamma, P) = \sup_{\eta > 0, \zeta \in \mathbb{R}, \lambda \in \Lambda} -\eta \mathbb{E}^{F_*}\left[\phi^\star\left(\frac{k(U,\theta,\gamma) + \zeta + \lambda' g(U,\theta,\gamma)}{-\eta}\right)\right] - \eta\delta - \zeta - \lambda'_{12}P, \quad (13)$$

$$\overline{K}_\delta(\theta; \gamma, P) = \inf_{\eta > 0, \zeta \in \mathbb{R}, \lambda \in \Lambda} \eta \mathbb{E}^{F_*}\left[\phi^\star\left(\frac{k(U,\theta,\gamma) - \zeta - \lambda' g(U,\theta,\gamma)}{\eta}\right)\right] + \eta\delta + \zeta + \lambda'_{12}P. \quad (14)$$

Moreover, the value of (13) is $+\infty$ (equivalently, the value of (14) is $-\infty$) if and only if there is no distribution in $\mathcal{N}_\delta$ under which (1) holds at (θ, γ, P).

Remark 2.6 *Problems (13) and (14) are convex in (η, ζ, λ). The parameter η is the Lagrange multiplier for the constraint $D_\phi(F\|F_*) \leq \delta$. Similarly, λ collects the Lagrange multipliers for the moment (in)equalities (1a)–(1d). These multipliers are non-negative if they correspond to inequality restrictions and unconstrained otherwise. Finally, ζ is the Lagrange multiplier for the constraint $\int \mathrm{d}F = 1$, which ensures that the optimization is over probability measures.*

Problems (13) and (14) simplify in some special cases. For KL neighborhoods, $\phi^\star(x) = e^x - 1$ and the multiplier ζ has a closed-form solution, leading to

$$\underline{K}_\delta(\theta; \gamma, P) = \sup_{\eta > 0, \lambda \in \Lambda} -\eta \log \mathbb{E}^{F_*} \left[e^{-(k(U, \theta, \gamma) + \lambda' g(U, \theta, \gamma))/\eta} \right] - \eta \delta - \lambda'_{12} P \,,$$

$$\overline{K}_\delta(\theta; \gamma, P) = \inf_{\eta > 0, \lambda \in \Lambda} \eta \log \mathbb{E}^{F_*} \left[e^{(k(U, \theta, \gamma) - \lambda' g(U, \theta, \gamma))/\eta} \right] + \eta \delta + \lambda'_{12} P \,.$$

Another special case is when $k(u, \theta, \gamma)$ does not depend on u. To analyze this case, consider

$$\Delta(\theta; \gamma, P) := \inf_F D_\phi(F \| F_*) \quad \text{subject to (1) holding at } (\theta, \gamma, P). \tag{15}$$

The value $\Delta(\theta; \gamma, P)$ is the minimum ϕ-divergence between F_* and a distribution F for which the moment conditions hold at (θ, γ, P). Proposition G.2 of Christensen and Connault (2022) shows that $\Delta(\theta; \gamma, P)$ has an equivalent dual formulation:

$$\Delta(\theta; \gamma, P) = \sup_{\zeta \in \mathbb{R}, \lambda \in \Lambda} -\mathbb{E}^{F_*} \left[\phi^\star(-\zeta - \lambda' g(U, \theta, \gamma)) \right] - \zeta - \lambda'_{12} P \,. \tag{16}$$

For KL divergence, ζ may be solved for in closed-form and problem (16) simplifies to

$$\Delta(\theta; \gamma, P) = \sup_{\lambda \in \Lambda} -\log \mathbb{E}^{F_*} \left[e^{-\lambda' g(U, \theta, \gamma)} \right] - \lambda'_{12} P \,.$$

When k does not depend on u, by a change of variables[11] we may then restate problems (13) and (14) as

$$\underline{K}_\delta(\theta; \gamma, P) = \begin{bmatrix} k(\theta, \gamma) \\ +\infty \end{bmatrix} \,, \quad \overline{K}_\delta(\theta; \gamma, P) = \begin{bmatrix} k(\theta, \gamma) & \text{if } \Delta(\theta; \gamma, P) \leq \delta, \\ -\infty & \text{if } \Delta(\theta; \gamma, P) > \delta. \end{bmatrix} \tag{17}$$

An important feature of our approach is that the optimization problems (13), (14), and (16) are convex and their dimension does not increase with δ. This feature is not shared by other seemingly natural approaches to flexibly model F, such as mixtures or other finite-dimensional sieves. As we show in Section 2.5, our procedure may be used to approximate sharp nonparametric bounds on counterfactuals by taking δ to be large but finite.

---

[11]Substitute $\eta\zeta - k(\theta, \gamma)$ in place of ζ in (13) and $\eta\zeta + k(\theta, \gamma)$ in place of ζ in (14), then substitute $\eta\lambda$ in place of λ in both (13) and (14).

2.4 Estimation

We now propose simple estimators of the bounds $\underline{\kappa}_\delta$ and $\overline{\kappa}_\delta$ based on "plugging in" consistent estimators $(\hat{P}, \hat{\gamma})$ of (P_0, γ_0). Estimators $\hat{\underline{\kappa}}_\delta$ and $\hat{\overline{\kappa}}_\delta$ are computed by optimizing criterion functions with respect to θ:

$$\hat{\underline{\kappa}}_\delta = \inf_{\theta \in \Theta} \hat{\underline{K}}_\delta(\theta), \qquad \hat{\overline{\kappa}}_\delta = \sup_{\theta \in \Theta} \hat{\overline{K}}_\delta(\theta),$$

where

$$\hat{\underline{K}}_\delta(\theta) = \left[\begin{array}{l} \underline{K}_\delta(\theta; \hat{\gamma}, \hat{P}) \\ +\infty \end{array} \right., \quad \hat{\overline{K}}_\delta(\theta) = \left[\begin{array}{ll} \overline{K}_\delta(\theta; \hat{\gamma}, \hat{P}) & \text{if } \Delta(\theta; \hat{\gamma}, \hat{P}) < \delta, \\ -\infty & \text{if } \Delta(\theta; \hat{\gamma}, \hat{P}) \geq \delta, \end{array} \right.$$

and $\underline{K}_\delta(\theta; \hat{\gamma}, \hat{P})$, $\overline{K}_\delta(\theta; \hat{\gamma}, \hat{P})$, and $\Delta(\theta; \hat{\gamma}, \hat{P})$ are the criterion functions (13), (14), and (16) evaluated at $(\hat{\gamma}, \hat{P})$. If $k(u, \theta, \gamma) = k(\theta, \gamma)$, then we simply have

$$\hat{\underline{K}}_\delta(\theta) = \left[\begin{array}{l} k(\theta, \hat{\gamma}) \\ +\infty \end{array} \right., \quad \hat{\overline{K}}_\delta(\theta) = \left[\begin{array}{ll} k(\theta, \hat{\gamma}) & \text{if } \Delta(\theta; \hat{\gamma}, \hat{P}) < \delta, \\ -\infty & \text{if } \Delta(\theta; \hat{\gamma}, \hat{P}) \geq \delta. \end{array} \right.$$

In Section 6.1 we establish consistency of $\hat{\underline{\kappa}}_\delta$ and $\hat{\overline{\kappa}}_\delta$ and derive their asymptotic distribution.

2.5 Nonparametric Bounds on Counterfactuals

We define the (nonparametric) identified set of counterfactuals as

$$\mathcal{K} = \left\{ \mathbb{E}^F[k(U, \theta, \gamma_0)] : (1) \text{ holds for some } \theta \in \Theta \text{ and } F \in \mathcal{F}_\theta \right\},$$

where $\mathcal{F}_\theta = \{F \in \mathcal{F} : \mathbb{E}^F[g(U, \theta, \gamma_0)] \text{ is finite and } F \ll \mu\}$ denotes all distributions on $\mathcal{U}$ that are absolutely continuous with respect to a σ-finite dominating measure μ and for which the moments in (1) are finite at θ. We impose existence of a density with respect to μ as it is often a structural assumption used, e.g., to avoid ties in CCPs or to establish existence of equilibria. The main result of this section shows that $\underline{\kappa}_\delta$ and $\overline{\kappa}_\delta$ approach the sharp nonparametric bounds $\inf \mathcal{K}$ and $\sup \mathcal{K}$ as δ becomes large.

We first introduce some additional regularity conditions. Say k is "μ-essentially bounded" if $|k(\cdot, \theta, \gamma_0)|$ has finite μ-essential supremum[12] for each $\theta \in \Theta$. This holds

---

[12]The μ-essential supremum of a function f is denoted μ-ess sup f and is the smallest value c for

trivially if k is bounded (e.g., counterfactual CCPs in Examples 2.2 and 2.3 and change in average welfare in Example 2.3). Models with unbounded k may be reparameterized (as a proof device) by setting $\tilde{\theta} = (\theta, \kappa)$, appending $k(U, \theta, \gamma_0) - \kappa$ as an element of g_4, and setting $k(U, \tilde{\theta}, \gamma_0) = \kappa$.

We also require a constraint qualification condition. This is a sufficient condition for establishing equivalence of "nonparametric" primal and dual problems in Appendix B, which is an intermediate step in the proof of the following result. Let 0_{d_i} denote a $d_i \times 1$ vector of zeros, $\mathcal{C} = \mathbb{R}_+^{d_1} \times \{0_{d_2}\} \times \mathbb{R}_+^{d_3} \times \{0_{d_4}\}$, $\mathcal{G}(\theta, \gamma) = \{\mathbb{E}^F[g(U, \theta, \gamma)] : F \in \mathcal{N}_\infty\}$ where $\mathcal{N}_\infty = \{F : D_\phi(F\|F_*) < \infty\}$, and $\vec{P} = (P, 0_{d_3+d_4})$. For $A, B \subseteq \mathbb{R}^d$, we let $\text{ri}(A)$ denote the relative interior of A and $A + B = \{a + b : a \in A, b \in B\}$.

Definition 2.1 Condition S *holds at* (θ, γ, P) *if* $\vec{P} \in \text{ri}(\mathcal{G}(\theta, \gamma) + \mathcal{C})$.

Using relative interior instead of interior allows for moment functions that are collinear at some θ (i.e., some moments are redundant). To give some intuition, consider moment equality models. Condition S requires that (1) holds at (θ, γ, P) under some $F \in \mathcal{N}_\delta$ that is "interior" to $\mathcal{N}_\infty$, in the sense that one can perturb the (non-redundant) moments in any direction by perturbing F. For moment inequality models, Condition S also requires that there is $F \in \mathcal{N}_\infty$ under which all moment inequalities hold strictly at (θ, γ, P).

Let $\Theta_I = \{\theta \in \Theta : (1) \text{ holds for some } F \in \mathcal{F}_\theta\}$ denote the (nonparametric) identified set for θ. Define the "nonparametric" objective function

$$\underline{K}_{np}(\theta; \gamma, P) = \inf_{F \in \mathcal{F}_\theta} \mathbb{E}^F[k(U, \theta, \gamma)] \quad \text{subject to (1) holding at } (\theta, \gamma, P), \qquad (18)$$

with the understanding that $\underline{K}_{np}(\theta; \gamma, P) = +\infty$ if the infimum runs over an empty set. Let $\overline{K}_{np}(\theta; \gamma, P)$ denote the analogous supremum. Evidently,

$$\inf \mathcal{K} = \inf_{\theta \in \Theta} \underline{K}_{np}(\theta; \gamma_0, P_0) \quad \text{and} \quad \sup \mathcal{K} = \sup_{\theta \in \Theta} \overline{K}_{np}(\theta; \gamma_0, P_0).$$

Definition 2.2 Θ_I *is S-regular if for all* $\epsilon > 0$ *there exist* $\underline{\theta}, \overline{\theta} \in \Theta_I$ *such that Condition S holds at* $(\underline{\theta}, \gamma_0, P_0)$ *and* $(\overline{\theta}, \gamma_0, P_0)$, $\underline{K}_{np}(\underline{\theta}; \gamma_0, P_0) < \inf \mathcal{K} + \epsilon$, *and* $\overline{K}_{np}(\overline{\theta}; \gamma_0, P_0) > \sup \mathcal{K} - \epsilon$.

Intuitively, S-regularity requires that the values the counterfactual takes at "boundary" points of Θ_I (i.e., at which Condition S fails) are not materially more extreme than values

---

which $\mu(\{u : f(u) > c\}) = 0$. The μ-essential infimum, denoted $\mu\text{-ess}\inf$, is defined analogously.

it can take at points "inside" Θ_I (i.e., at which Condition S holds). This condition can be verified under more primitive continuity conditions on k and g. A sufficient (but not necessary) condition for S-regularity is that Condition S holds at (θ, γ_0, P_0) for all $\theta \in \Theta_I$.

Theorem 2.1 *Suppose that Assumption Φ holds, k is μ-essentially bounded, Θ_I is S-regular, and μ and F_* are mutually absolutely continuous. Then*

$$\lim_{\delta \to \infty} \underline{\kappa}_\delta = \inf \mathcal{K}, \qquad\qquad \lim_{\delta \to \infty} \overline{\kappa}_\delta = \sup \mathcal{K}.$$

Theorem 2.1 shows that our procedure can be used to approximate the sharp non-parametric bounds $\inf \mathcal{K}$ and $\sup \mathcal{K}$ by setting δ to be large but finite. If μ is Lebesgue measure—which it often is in applications—then the mutual absolute continuity condition in Theorem 2.1 is satisfied whenever F_* has strictly positive density over $\mathcal{U}$.

Remark 2.7 *Appendix B presents the dual forms of $\underline{K}_{np}$ and $\overline{K}_{np}$. Unlike $\underline{K}_\delta$ and $\overline{K}_\delta$, the duals of $\underline{K}_{np}$ and $\overline{K}_{np}$ are min-max and max-min problems which involve an inner optimization over u. These problems may be computationally challenging, especially when u is multivariate. Comparing Proposition 2.1 with the duals in Appendix B, we see that setting $\delta < \infty$ replaces a "hard-max" (an optimization over u) with a "soft-max" (a convex expectation). In this respect, adding the constraint $F \in \mathcal{N}_\delta$ may be viewed as a regularization of the nonparametric objective functions, similar to the use of entropic penalization to regularize objective functions in optimal transport problems—see, e.g., Cuturi (2013). Smaller values of δ impose a stronger regularization.*

Theorem 2.1 is silent on the issue of how large δ needs to be so that $\underline{\kappa}_\delta$ and $\overline{\kappa}_\delta$ are close to the nonparametric bounds. While this is model- and counterfactual-specific, the following toy example suggests that relatively small values of δ may suffice in some problems where the counterfactual is a choice probability.

Example 2.4 Consider the problem

$$\overline{\kappa}_\delta = \sup_{\theta \in \mathbb{R}, F \in \mathcal{N}_\delta} \mathbb{E}^F[\mathbb{1}\{U \leq \theta\}] \quad \text{subject to} \quad \mathbb{E}^F[U - \theta] = 0,$$

where $\mathcal{N}_\delta$ is defined by KL divergence and F_* is the $N(0,1)$ distribution. When $F = F_*$, the only solution to $\mathbb{E}^F[U - \theta] = 0$ is $\theta = 0$. Therefore, the value of the counterfactual under F_* is $\mathbb{E}^{F_*}[\mathbb{1}\{U \leq 0\}] = \frac{1}{2}$ whereas $\sup \mathcal{K} = 1$. The large-δ approximation $\overline{\kappa}_\delta =$

16

$1 - 2\pi e^{-2\delta-1}(1 + o(1))$ is derived in Appendix H of Christensen and Connault (2022). By symmetry, $\underline{\kappa}_\delta = 2\pi e^{-2\delta-1}(1+o(1))$ and $\inf \mathcal{K} = 0$. Therefore, in this example, $\underline{\kappa}_\delta$ and $\overline{\kappa}_\delta$ converge rapidly to $\inf \mathcal{K}$ and $\sup \mathcal{K}$ as δ increases. $\qquad\square$

More generally, suppose the dual problems (13) and (14) have unique solutions $\underline{\eta}$ and $\overline{\eta}$ for η, where the optimization is performed over $\eta \geq 0$.[13] Under appropriate regularity conditions (see, e.g., Milgrom and Segal (2002)), it follows that

$$\frac{\partial \underline{K}_\delta(\theta; \gamma, P)}{\partial \delta} = -\underline{\eta}, \qquad \frac{\partial \overline{K}_\delta(\theta; \gamma, P)}{\partial \delta} = \overline{\eta}.$$

One can therefore infer from $\underline{\eta}$ and $\overline{\eta}$ the extent to which, if at all, the bounds at any fixed θ would widen further if δ was increased.

3   Practical Considerations

We now discuss practical details for implementing our procedure. Section 3.1 discusses computational methods, Section 3.2 presents our MPEC approach, and Section 3.3 discusses methods for dealing with over-identified models.

3.1   Computation

There are three aspects to computation: (i) computing the expectations with respect to F_* in the objective functions, (ii) solving the inner optimization problems over Lagrange multipliers, and (iii) solving the outer optimization problems over θ.

The expectations in the objective functions (13), (14), and (16) are available in closed form for certain settings,[14] in which case the dimension of u does not play a role in the computational complexity of our procedure. Otherwise, the expectations will need to be computed numerically. If so, the dimension of u will play a role in terms of determining how many quadrature points or Monte Carlo draws are needed to control numerical approximation error. In the empirical applications we used a randomized quasi-Monte Carlo approach based on scrambled Halton sequences as in Owen (2017).

---

[13] Optimizing over $\eta \geq 0$ rather than $\eta > 0$ does not affect the optimal value—see Proposition G.1 of Christensen and Connault (2022).

[14] An earlier draft derived closed-form expressions for a discrete game of complete information with Gaussian payoff shocks and KL neighborhoods—see https://arxiv.org/abs/1904.00989v2.

The inner optimization with respect to Lagrange multipliers can be solved rapidly: it is convex and gradients and Hessians are available in closed-form. The envelope theorem can be used to derive gradients for the outer optimization when k and g are differentiable in θ.[15] Our procedures were all implemented in Julia with the inner and outer optimizations solved using Knitro. A general-purpose implementation of our methods in Julia is provided in the supplemental material.

As with parameter estimation in nonlinear structural models, the outer optimization with respect to θ is typically non-convex. In applications, we used an iterative multi-start procedure in an attempt to converge to global optima. Computation times are reported in the applications below.

3.2 MPEC Approach

We now describe and formally justify an MPEC version of our procedure in the spirit of Su and Judd (2012). This approach simplifies computation in models with endogenous parameters defined by equilibrium conditions (e.g., value functions defined by Bellman equations), resulting in significant computational gains for DDC models in particular.

Suppose $\theta = (\theta_s, \theta_e)$ and $g_4 = (g_{4s}, g_{4e})$ where θ_s are "deep" structural parameters and θ_e are "endogenous" parameters that are defined implicitly by g_{4e}. That is, for any (θ_s, γ, F), the parameter $\theta_e = \theta_e(\theta_s, \gamma, F)$ solves

$$\mathbb{E}^F[g_{4e}(U, (\theta_s, \theta_e), \gamma)] = 0 \,.$$

For instance, in Example 2.3 we have $\theta_s = \theta_\pi$ or (θ_π, β), while $\theta_e = (v, \tilde{v})$ collects the value functions in the baseline model and counterfactual, and g_{4e} collects the functions representing the corresponding Bellman equations, as in display (6). Although our procedure can be implemented as described in Section 2, that implementation does not make use of the fact that θ_e is defined implicitly by g_{4e}.

To leverage this structure, consider the subset of moments conditions excluding g_{4e}:

$$\begin{aligned}
\mathbb{E}^F[g_1(U, \theta, \gamma_0)] &\le P_{10}, \quad \mathbb{E}^F[g_2(U, \theta, \gamma_0)] = P_{20}, \\
\mathbb{E}^F[g_3(U, \theta, \gamma_0)] &\le 0, \qquad \mathbb{E}^F[g_{4s}(U, \theta, \gamma_0)] = 0,
\end{aligned} \tag{19}$$

---

[15]In practice, we smoothed any non-smooth moments and used automatic differentiation to compute derivatives with respect to θ if these were not easily available analytically.

and define criterion functions using these only:

$$\underline{K}_\delta^s(\theta; \gamma, P) = \inf_{F \in \mathcal{N}_\delta} \mathbb{E}^F[k(U, \theta, \gamma)] \quad \text{subject to (19) holding at } (\theta, \gamma, P), \qquad (20)$$

$$\overline{K}_\delta^s(\theta; \gamma, P) = \sup_{F \in \mathcal{N}_\delta} \mathbb{E}^F[k(U, \theta, \gamma)] \quad \text{subject to (19) holding at } (\theta, \gamma, P). \qquad (21)$$

Under the conditions of Proposition 2.1, these criterion functions may be restated as

$$\underline{K}_\delta^s(\theta; \gamma, P) = \sup_{\eta>0, \zeta \in \mathbb{R}, \lambda \in \Lambda_s} -\eta \mathbb{E}^{F_*} \left[\phi^\star \left(\frac{k(U, \theta, \gamma) + \zeta + \lambda' g_s(U, \theta, \gamma)}{-\eta} \right) \right] - \eta\delta - \zeta - \lambda_{12}' P, \qquad (22)$$

$$\overline{K}_\delta^s(\theta; \gamma, P) = \inf_{\eta>0, \zeta \in \mathbb{R}, \lambda \in \Lambda_s} \eta \mathbb{E}^{F_*} \left[\phi^\star \left(\frac{k(U, \theta, \gamma) - \zeta - \lambda' g_s(U, \theta, \gamma)}{\eta} \right) \right] + \eta\delta + \zeta + \lambda_{12}' P, \qquad (23)$$

with $g_s = (g_1, g_2, g_3, g_{4s})$ and $\Lambda_s = \mathbb{R}_+^{d_1} \times \mathbb{R}^{d_2} \times \mathbb{R}_+^{d_3} \times \mathbb{R}^{d_{4s}}$ with $d_{4s} = \dim(g_{4s})$. Problems (22) and (23) simplify analogously to (17) when k does not depend on u, with the minimum divergence problem Δ defined using g_s in place of g.

In our MPEC approach, the criterion functions (22) and (23) are optimized with respect to θ, with the remaining moment conditions involving g_{4e} appended as constraints. Importantly, these constraints are evaluated under the "least favorable" distributions $\underline{F}_{\delta,\theta}$ and $\overline{F}_{\delta,\theta}$ that solve problems (20) and (21), respectively. The following proposition formally justifies this approach.

Proposition 3.1 *Suppose that Assumption Φ holds. Then the problems*

$$\inf_{\theta \in \Theta} \underline{K}_\delta(\theta; \gamma, P)$$

and

$$\inf_{\theta \in \Theta} \underline{K}_\delta^s(\theta; \gamma, P) \quad \text{subject to} \quad \mathbb{E}^{\underline{F}_{\delta,\theta}}[g_{4e}(U, \theta, \gamma)] = 0$$

have the same value. An analogous result holds for the upper bound.

To implement our MPEC approach, note that the expectations in the constraints may be expressed in terms of changes of measure. Let $\underline{m}_{\delta,\theta} = \mathrm{d}\underline{F}_{\delta,\theta}/\mathrm{d}F_*$ and $\overline{m}_{\delta,\theta} = \mathrm{d}\overline{F}_{\delta,\theta}/\mathrm{d}F_*$ so that

$$\mathbb{E}^{\underline{F}_{\delta,\theta}}[\,\cdot\,] = \mathbb{E}^{F_*}[\underline{m}_{\delta,\theta}(U) \,\cdot\,], \quad \mathbb{E}^{\overline{F}_{\delta,\theta}}[\,\cdot\,] = \mathbb{E}^{F_*}[\overline{m}_{\delta,\theta}(U) \,\cdot\,].$$

If k depends on u, then we construct $\underline{m}_{\delta,\theta}$ and $\overline{m}_{\delta,\theta}$ from solutions to (22) and (23), say $(\underline{\eta}, \underline{\zeta}, \underline{\lambda})$ and $(\overline{\eta}, \overline{\zeta}, \overline{\lambda})$ (these solutions exist under the regularity conditions below). If $\underline{\eta} > 0$, then the distribution solving (20) is unique and is induced by the change of

measure

$$\underline{m}_{\delta,\theta}(u) = \dot{\phi}^\star\left(\frac{k(u,\theta,\gamma) + \underline{\zeta} + \underline{\lambda}'g_s(u,\theta,\gamma)}{-\underline{\eta}}\right), \tag{24}$$

where $\dot{\phi}^\star(x) = \frac{d\phi^\star(x)}{dx}$. The function $\overline{m}_{\delta,\theta}(u)$ is constructed similarly, replacing $(\underline{\eta}, \underline{\zeta}, \underline{\lambda})$ in (24) by $(-\overline{\eta}, -\overline{\zeta}, -\overline{\lambda})$. For KL divergence the change of measure simplifies to

$$\underline{m}_{\delta,\theta}(u) = \frac{e^{(k(u,\theta,\gamma)+\underline{\lambda}'g_s(u,\theta,\gamma))/-\underline{\eta}}}{\mathbb{E}^{F_*}\left[e^{(k(u,\theta,\gamma)+\underline{\lambda}'g_s(u,\theta,\gamma))/-\underline{\eta}}\right]},$$

and similarly for $\underline{m}_{\delta,\theta}(u)$.

If $\underline{\eta} = 0$, then there may be multiple minimizing distributions. As shown in the proof of Proposition 3.2, each such distribution must be supported on

$$\underline{A}_{\delta,\theta} := \{u : k(u,\theta,\gamma) + \underline{\lambda}'g_s(u,\theta,\gamma) = F_*\text{-ess}\inf(k(\cdot,\theta,\gamma) + \underline{\lambda}'g_s(\cdot,\theta,\gamma))\}.$$

Note $F_*(\underline{A}_{\delta,\theta}) > 0$ is required for $\underline{\eta} = 0$ to be a solution. Otherwise, any distribution supported on $\underline{A}_{\delta,\theta}$ is not absolutely continuous with respect to F_* and is therefore not in $\mathcal{N}_\delta$. If $\underline{\eta} = 0$ and $F_*(\underline{A}_{\delta,\theta}) > 0$, then we construct $\underline{m}_{\delta,\theta}$ by restricting F_* to $\underline{A}_{\delta,\theta}$ and rescaling:

$$\underline{m}_{\delta,\theta}(u) = \mathbb{1}\{u \in \underline{A}_{\delta,\theta}\}/F_*(\underline{A}_{\delta,\theta}).$$

The function $\overline{m}_{\delta,\theta}(u)$ is constructed analogously, replacing $\underline{\lambda}$ with $-\overline{\lambda}$ and the set $\underline{A}_{\delta,\theta}$ with $\overline{A}_{\delta,\theta} = \{u : k(u,\theta,\gamma) - \overline{\lambda}'g_s(u,\theta,\gamma) = F_*\text{-ess}\sup(k(\cdot,\theta,\gamma) - \overline{\lambda}'g_s(\cdot,\theta,\gamma))\}$.

If k does not depend on u, then $\underline{m}_{\delta,\theta}$ and $\overline{m}_{\delta,\theta}$ are constructed from solutions to a version of problem (16) with g_s in place of g. Under the regularity conditions below, this program has a solution, say $(\underline{\zeta}, \underline{\lambda})$. In this case, we define

$$\underline{m}_{\delta,\theta}(u) = \overline{m}_{\delta,\theta}(u) = \dot{\phi}^\star\left(-\underline{\zeta} - \underline{\lambda}'g_s(u,\theta,\gamma)\right). \tag{25}$$

For KL divergence the change of measure simplifies to

$$\underline{m}_{\delta,\theta}(u) = \overline{m}_{\delta,\theta}(u) = \frac{e^{-\underline{\lambda}'g_s(u,\theta,\gamma)}}{\mathbb{E}^{F_*}\left[e^{-\underline{\lambda}'g_s(u,\theta,\gamma)}\right]}.$$

Proposition 3.2 *Suppose that Assumption Φ holds, Condition S holds at (θ, γ, P), and there exists a distribution F with $D(F\|F_*) < \delta$ under which (19) holds at (θ, γ, P). Then the distributions $\underline{F}_{\delta,\theta}$ and $\overline{F}_{\delta,\theta}$ induced by $\underline{m}_{\delta,\theta}$ and $\overline{m}_{\delta,\theta}$ solve (20) and (21), respectively.*

Table 2: Computation times (in seconds) for the inner problems

Implementation	Objective			
	$\overline{K}_{0.01}$	$\overline{K}_{0.10}$	$\overline{K}_{1.00}$	Δ
MPEC (92 moments)	0.207	0.232	0.256	0.108
Full (272 moments)	4.317	12.978	43.699	3.365

Note: Expectations are computed using 50,000 scrambled Halton draws. Computations are performed in Julia v1.6.4 and Knitro v12.4.0 on a 2.7GHz MacBook Pro with 16GB memory.

Example. We consider a numerical example for the DDC model of Rust (1987) based on the parameterization in Section 5.4 of Norets and Tang (2014). The counterfactual they consider is a hypothetical change in the law of motion of the state. We follow these papers and use state-space of dimension 90. As $|\mathcal{S}| = 90$ and $\mathcal{D}_0 = \{0, 1\}$, there are 90 functions in g_2 representing the observed CCPs. There are another 180 functions in g_{4e} representing the Bellman equations in the baseline model and counterfactual across states. We also impose the normalization $\mathbb{E}^F[U_d] = 0$ for $d = 0, 1$. Hence, $g_{4s}(U, \theta, \gamma) = (U_0, U_1)$. Our MPEC approach has 92 moments in the inner optimization (90 for CCPs and two mean-zero normalizations on the shocks) with the remaining 180 moments representing the Bellman equations appended as constraints. The full approach uses all 272 moments in the inner optimization.

Table 2 reports computation times for the inner optimization problems (14) and (23) (denoted $\overline{K}_\delta$) for maximizing the counterfactual CCP in the highest mileage state.[16] We also report times for solving the minimum divergence problem (16) (denoted Δ) using the full set of moment functions g and its MPEC analogue using g_s. Neighborhoods are constrained by a hybrid of KL and χ^2 divergence as in the empirical applications—see Section 5. As can be seen, the inner optimization problems are solved at least 20 times faster for the MPEC implementation, with the relative efficiency increasing in δ.

3.3 Over-identification

In over-identified models (i.e., where the number of moment conditions d exceeds the dimension d_θ of θ), there might not exist $\theta \in \Theta$ for which the sample moment conditions (7) hold under F_*. We propose two methods for handling over-identified models.

---

[16]The times in Table 2 are based on initializing the solver at $\eta = 1$, $\zeta = 0$, and $\lambda = 0$. When embedded in the outer optimization over θ, computation times for the inner problem are reduced significantly by using a warm start that initializes at the solution to the inner problem at the previous value of θ.

First, one may compute the smallest value of δ for which there exists $F \in \mathcal{N}_\delta$ consistent with the sample moment conditions (7) by solving the optimization problem

$$\hat{\delta} = \inf_{\theta \in \Theta} \Delta(\theta; \hat{\gamma}, \hat{P}).$$

The interval $[\hat{\underline{\kappa}}_\delta, \hat{\overline{\kappa}}_\delta]$ will be nonempty for $\delta > \hat{\delta}$. If the model is correctly specified under F_*,[17] then $\hat{\delta}$ will converge in probability to zero under the conditions of Theorem 6.1. In this case, the interval $[\hat{\underline{\kappa}}_\delta, \hat{\overline{\kappa}}_\delta]$ will be nonempty with probability approaching one for each fixed $\delta > 0$.

It is also possible that $\hat{\delta} = +\infty$ in correctly specified but over-identified models when $\hat{P}$ is incompatible with certain model restrictions. For instance, CCPs are often estimated nonparametrically using empirical choice frequencies. If some choices aren't observed in the data, then the estimated CCPs will be zero even though model-implied CCPs are strictly positive.

This issue can be circumvented in models defined by equality restrictions only (hence $P_0 \equiv P_{20}$) using the following two-step approach. First, compute a preliminary estimator $\tilde{\theta}$ of θ based on (7). Then, set $\hat{P} = \mathbb{E}^{F_*}[g_2(U, \tilde{\theta}, \hat{\gamma})]$. This second-step estimator $\hat{P}$ is compatible with the model by construction, thereby ensuring that the interval $[\hat{\underline{\kappa}}_\delta, \hat{\overline{\kappa}}_\delta]$ is nonempty for each $\delta > 0$. The estimator $\hat{P}$ will be consistent and asymptotically normal under mild regularity conditions provided the model is correctly specified under F_*, so the consistency and inference results developed in Section 6 will also apply.

4 Interpreting the Neighborhood Size

This section presents some theoretical results and practical methods to help interpret the neighborhood size δ. Sections 4.1 and 4.4 discuss properties of ϕ-divergences and their implications for interpreting δ. Section 4.2 shows how to construct the "least favorable" distributions that minimize or maximize the counterfactual. Section 4.3 gives a practical, model-based metric for interpreting δ.

---

[17]Neither our theoretical results developed in Section 2 nor the estimation and inference results in Section 6 require correct specification of the model under F_*.

4.1 Invariance

A defining property of ϕ-divergences are their invariance to invertible transformations. That is, if T is an invertible transformation and G and G_* denote the distributions of $T(U)$ when $U \sim F$ and $U \sim F_*$, respectively, then $D_\phi(F\|F_*) = D_\phi(G\|G_*)$.[18] An important consequence of invariance is that δ has the same interpretation under a change in units. For instance, if one researcher writes a model in terms of dollars with $U \sim F_*$ and another researcher uses thousands of dollars with $U \sim G_*$ for $G_*(u) = F_*(10^{-3}u)$, then F is in $\mathcal{N}_\delta$ if and only if its rescaled counterpart G is in a δ-neighborhood of G_*. A second consequence is that neighborhood size is invariant under invertible location and scale transformations of F_* (e.g., $N(\mu, \Sigma)$ versus $N(0, I)$).

4.2 Least Favorable Distributions

A useful feature of our approach is that the "least favorable" distributions (LFDs) that attain the smallest or largest values of the counterfactual may easily be recovered. To help interpret δ, one may plot the LFDs and compute other quantities of interest (e.g., correlations or welfare measures) under them.

Section 3.2 describes how to construct LFDs when our MPEC approach is used. LFDs for our full (i.e., non-MPEC) approach are a special case with $g_4 = g_{4s}$. To briefly summarize, consider the LFD $\underline{F}_{\delta,\theta}$ solving the minimization problem (11). First suppose that k depends on u. Let $(\underline{\eta}, \underline{\zeta}, \underline{\lambda})$ solve problem (13). If $\underline{\eta} > 0$, then $\underline{F}_{\delta,\theta}$ is unique and its change-of-measure $\underline{m}_{\delta,\theta} = \mathrm{d}\underline{F}_{\delta,\theta}/\mathrm{d}F_*$ is given by

$$\underline{m}_{\delta,\theta}(u) = \dot{\phi}^\star \left(\frac{k(u,\theta,\gamma) + \underline{\zeta} + \underline{\lambda}' g(u,\theta,\gamma)}{-\underline{\eta}} \right). \tag{26}$$

The LFD $\overline{F}_{\delta,\theta}$ solving the maximization problem (12) is constructed similarly, replacing $(\underline{\eta}, \underline{\zeta}, \underline{\lambda})$ in (26) with $(-\overline{\eta}, -\overline{\zeta}, -\overline{\lambda})$, where $(\overline{\eta}, \overline{\zeta}, \overline{\lambda})$ solves (14). If $\underline{\eta} = 0$ or $\overline{\eta} = 0$, then there may exist multiple distributions solving (11) and (12) at θ. LFDs in this case are constructed analogously to the method described in Section 3.2. Note that $\underline{\eta} = 0$ or $\overline{\eta} = 0$ is unlikely if k and/or elements of g are unbounded in u—see the discussion in Section 3.2. If k does not depend on u, then we set

$$\underline{m}_{\delta,\theta}(u) = \overline{m}_{\delta,\theta}(u) = \dot{\phi}^\star \left(-\underline{\zeta} - \underline{\lambda}' g(u,\theta,\gamma) \right) \tag{27}$$

---

[18]See, e.g., Liese and Vajda (1987). A more direct statement is in Qiao and Minematsu (2010).

where $(\underline{\zeta}, \underline{\lambda})$ solves (16). While there may exist multiple distributions solving (11) and (12) in this case, the distribution induced by (27) has smallest ϕ-divergence relative to F_* among all such distributions.

4.3 Viewing Neighborhood Size through the Lens of the Model

Another method for interpreting δ is based on measuring the variation in the moments at the distributions solving (8) and (9) relative to their values under F_*.

Consider the sets of minimizing and maximizing values of θ at which $\underline{\kappa}_\delta$ and $\overline{\kappa}_\delta$ are attained, say $\underline{\Theta}_\delta$ and $\overline{\Theta}_\delta$. These are nonempty under the regularity conditions in Section 6. While the moment conditions (1) hold at any $\theta \in \underline{\Theta}_\delta \cup \overline{\Theta}_\delta$ under the corresponding LFD, they will typically not hold at θ under F_*. We therefore define

$$size(\delta) = \sup_{\theta \in \underline{\Theta}_\delta \cup \overline{\Theta}_\delta} \max \left\{ \left\| \left(\mathbb{E}^{F_*}[g_1(U, \theta, \gamma_0)] - P_{10} \right)_+ \right\|_\infty, \left\| \mathbb{E}^{F_*}[g_1(U, \theta, \gamma_0)] - P_{20} \right\|_\infty, \right.$$
$$\left. \left\| \left(\mathbb{E}^{F_*}[g_3(U, \theta, \gamma_0)] \right)_+ \right\|_\infty, \left\| \mathbb{E}^{F_*}[g_4(U, \theta, \gamma_0)] \right\|_\infty \right\},$$

where $(v)_+ = (\max\{v_i, 0\})_{i=1}^d$ for a vector $v \in \mathbb{R}^d$. The quantity $size(\delta)$ is the maximum degree to which the moments at $\theta \in \underline{\Theta}_\delta \cup \overline{\Theta}_\delta$ violate (1) under F_*.

This measure is informative about the extent to which the distortions to F_* required to attain the smallest and largest values of the counterfactual over $\mathcal{N}_\delta$ are reflected in (1). Small values of $size(\delta)$ indicate that the LFDs supporting $\underline{\kappa}_\delta$ and $\overline{\kappa}_\delta$ distort F_* in a way that moves the counterfactual but barely moves the moments. Conversely, large values of $size(\delta)$ indicate that distortions required to increase or decrease the counterfactual also have a material impact on the moments. In practice, this measure can be computed by replacing (P_0, γ_0) by estimators $(\hat{P}, \hat{\gamma})$ and $\underline{\Theta}_\delta$ and $\overline{\Theta}_\delta$ by the minimizers and maximizers of the sample criterions or by the estimators of $\underline{\Theta}_\delta$ and $\overline{\Theta}_\delta$ introduced in Section 6.2.

4.4 Relating Different Divergences

It is well known that ϕ-divergences are equivalent over local neighborhoods (see, e.g., Theorem 4.1 of Csiszár and Shields (2004)). However, $\underline{\kappa}_\delta$ and $\overline{\kappa}_\delta$ may depend on the choice of ϕ when δ is not arbitrarily small. Bounds induced by different ϕ functions may be related as follows. Let $\mathcal{N}_{\delta,1}$ and $\mathcal{N}_{\delta,2}$ denote δ-neighborhoods induced by ϕ_1 and ϕ_2, respectively. The quantity

$$\bar{a} = \sup_{x \geq 0, x \neq 1} \frac{\phi_1(x)}{\phi_2(x)}$$

24

is a measure of relative neighborhood size: if $\bar{a} < \infty$ then $\mathcal{N}_{\delta,2} \subseteq \mathcal{N}_{\bar{a}\delta,1}$ for each $\delta > 0$, as shown formally in the proof of Proposition 4.1 below. For instance, when comparing KL divergence ($\phi_1(x) = x \log x - x + 1$) and χ^2 divergence ($\phi_2(x) = \frac{1}{2}(x-1)^2$) we obtain $\bar{a} = 2$. Therefore, δ-neighborhoods under χ^2 divergence are contained in 2δ-neighborhoods under KL divergence. Interchanging ϕ_1 and ϕ_2 produces $\bar{a} = +\infty$, which reflects the fact that KL divergence is weaker than χ^2 divergence.

Let $\underline{\kappa}_{\delta,1}$ and $\underline{\kappa}_{\delta,2}$ denote the smallest counterfactual from display (8) over $\mathcal{N}_{\delta,1}$ and $\mathcal{N}_{\delta,2}$, respectively. Define $\overline{\kappa}_{\delta,1}$ and $\overline{\kappa}_{\delta,2}$ analogously.

Proposition 4.1 *Suppose that Assumption Φ holds for both ϕ_1 and ϕ_2 and $\bar{a}$ is finite. Then $[\underline{\kappa}_{\delta,2}, \overline{\kappa}_{\delta,2}] \subseteq [\underline{\kappa}_{\bar{a}\delta,1}, \overline{\kappa}_{\bar{a}\delta,1}]$ for each $\delta > 0$.*

It follows from Proposition 4.1 that bounds that are wide under ϕ_2 must necessarily be wide under ϕ_1. Similarly, narrow bounds under ϕ_1 must also be narrow under ϕ_2. Note also that the inclusion in Proposition 4.1 holds for any counterfactual.

5 Empirical Applications

5.1 Marital College Premium

Chiappori et al. (2017), henceforth CSW, study the evolution of marital returns to education using a frictionless matching model with transferable utility (Choo and Siow, 2006). Within this framework, the "marital college premium" is the additional expected utility that an individual would derive from the marriage market if they had a (counterfactually) higher level of education. CSW find that marital college premiums for women in the United States increased significantly across cohorts from the mid to late 20th century, particularly for the more highly educated.

As is conventional following Dagsvik (2000) and Choo and Siow (2006), CSW assume latent variables representing individuals' idiosyncratic marital preferences are i.i.d. Gumbel. The marital college premium is only partially identified when the distribution of these latent variables is not specified. We therefore perform a sensitivity analysis of CSW's estimates to departures from this conventional parametric assumption.

Our analysis makes several findings. First, it seems impossible to draw conclusions about whether marital college premiums have increased or decreased over time under small nonparametric relaxations of the i.i.d. Gumbel assumption. Interestingly, premiums have narrow nonparametric bounds at fixed parameter values, but a slight relaxation

of the i.i.d. Gumbel assumption allows for significant variation in parameters which, in turn, produces uninformatively wide bounds. As parameters are just-identified under any fixed distribution of shocks (Galichon and Salanié, 2021), further restrictions on parameters or shape restrictions on the distribution are required to tighten the bounds. We show that imposing exchangeability can tighten the bounds significantly.

Model and Benchmark Estimates. Agents are male or female and one of J types (education levels). A type-a male receives utility ε_{a0} if he chooses to be unmatched and $z_{ab} + \varepsilon_{ab}$ if he matches with a type-b female. Similarly, a type-b female receives utility e_{0b} if she chooses to be unmatched and $t_{ab} + e_{ab}$ if she matches with a type-a male. The parameters $(z_{ab}, t_{ab})_{a,b=1}^{J}$ represent the common deterministic component of marital preferences. The latent shocks $(\varepsilon_{a0}, \ldots, \varepsilon_{aJ})$ and $(e_{0b}, \ldots, e_{Jb})$ represent individuals' idiosyncratic marital preferences. Shocks are i.i.d. across individuals and have mean zero. The type b to b' marital education premium for females is the difference in expected marital utility between types b and b':

$$\kappa = \mathbb{E}^{F} \left[\max_{a=0,\ldots,J} \left(t_{ab'} + e_{ab'} \right) \right] - \mathbb{E}^{F} \left[\max_{a=0,\ldots,J} \left(t_{ab} + e_{ab} \right) \right], \tag{28}$$

where F denotes the distribution of $(e_{0b}, \ldots, e_{Jb'})$ and $t_{0b} = t_{0b'} = 0$.

CSW use data from the American Community Survey. They form 28 cohorts indexed by female birth year from 1941 (cohort 1) to 1968 (cohort 28), each of which is treated as an independent marriage market. We focus on CSW's estimates for whites. There are $J = 5$ types: "high-school dropouts", "high-school graduates", "some college", "college graduate", and "college-plus". We center our analysis on the "some college" to "college graduate" premium, though we obtained qualitatively similar results (not reported) for the "college graduate" to "college-plus" premium. Figure 1 presents estimates and 95% confidence sets (CSs) for the premium under the i.i.d. Gumbel assumption (cf. Figure 21 in CSW) based on CSW's replication files.

Implementation. The model reduces to a standard individual-level discrete choice problem for each type (see CSW's Propositions 1 and 2). We assume that the distribution of females' preference shocks does not depend on their type, so we drop the b subscript and consider a single random vector $U = (e_0, \ldots, e_J)$. We allow the distribution F of U to vary across cohorts and implement our procedures cohort-by-cohort.[19]

---

[19]In view of the just-identification results of Galichon and Salanié (2021), we would obtain the same bounds if F was homogeneous across cohorts. Allowing for heterogeneity in own-type would result in wider bounds.

Under any fixed F, a cohort's parameters $(t_{ab})_{a=1}^{J}$ are just-identified from the marriage probabilities for that cohort's type-b women (Galichon and Salanié, 2021). We therefore impose only the moment conditions involving the parameters $\theta = (t_{ab}, t_{ab'})_{a=1}^{J}$ appearing in (28), as the remaining parameters can be chosen to fit the remaining marriage probabilities under the resulting least-favorable distribution. We form g_2 to explain the type b and b' marriage probabilities for women in a given cohort:

$$g_2(U, \theta) = \left[\begin{array}{c} (\mathbb{1}\{t_{ab} + e_a = \max_{a'=0,...,J}(t_{a'b} + e_{a'})\})_{a=1}^{J} \\ (\mathbb{1}\{t_{ab'} + e_a = \max_{a'=0,...,J}(t_{a'b'} + e_{a'})\})_{a=1}^{J} \end{array} \right]$$

and form $\hat{P}_2$ using CSW's estimates of the corresponding type-b and b' marriage probabilities. We set $g_4(U, \theta) = (e_j, e_j^2 - \pi^2/6)_{j=0}^{J}$ so that shocks have mean zero and the same variance as the Gumbel distribution. The scale normalization also ensures that the nonparametric bounds on the premium are finite at any fixed θ. As $J = 5$, there are 22 moments (10 for marriage probabilities and 12 location/scale normalizations), and θ has dimension 10.

We consider a second implementation which imposes invariance of F under rotations and reflections of potential spouse types, so that the model-implied marriage probabilities depend on θ but not the labeling of potential spouse types (though they may depend on their ordering).[20] Formally, this shape restriction corresponds to dihedral exchangeability (see Appendix A.1); we refer to it simply as "exchangeability". Under this shape restriction, F must satisfy the 22 moment conditions under all 12 rotations and reflections of the elements of U. This implementation therefore imposes a total of 264 moment conditions. Rather than including all 264 moments separately, it suffices to form g_2 and g_4 by taking the averages of the 22 moments across the 12 permutations (see Appendix A.1). Both implementations therefore have inner optimization problems of the same dimension.

Computations are performed as described in Section 3.1. The first implementation uses 50,000 scrambled Halton draws to compute the expectations. The second uses 10,000 draws which are concatenated over the 12 permutations (see Remark A.2), for a total of 120,000 draws. Computation times are reported in Appendix D.1. CSs for $\underline{\kappa}_\delta$ and $\overline{\kappa}_\delta$ are computed using the bootstrap procedure in Section 6.2. Appendix D.1 discusses bootstrap details and presents projection CSs using the method from Section 6.3.

---

[20]Allowing dependence on the ordering of types seems desirable here as types correspond to education levels, which are naturally ordered.

Table 3: Metrics for interpreting δ

	Without exchangeability			With exchangeability		
δ	$\rho_{\max}, \underline{\kappa}_\delta$	$\rho_{\max}, \overline{\kappa}_\delta$	$size$	$\rho_{\max}, \underline{\kappa}_\delta$	$\rho_{\max}, \overline{\kappa}_\delta$	$size$
0.01	-0.015	-0.014	0.010	-0.022	0.013	0.006
0.10	-0.071	-0.073	0.038	-0.061	0.054	0.023
1	-0.247	-0.197	0.112	-0.139	0.115	0.099
10	-0.502	-0.496	0.242	-0.204	0.236	0.176
100	-0.620	-0.576	0.266	-0.266	0.284	0.178

Note: Averages across cohorts of the largest element of the correlation matrix for U under the LFDs at which the estimated lower bounds ($\rho_{\max}$, $\underline{\kappa}_\delta$) and upper bounds ($\rho_{\max}$, $\overline{\kappa}_\delta$) are attained, and our *size* measure from Section 4.3. Each is computed at the parameter values at which the estimated upper and lower bounds are attained.

We define neighborhoods using a hybrid of KL and χ^2 divergence:

$$\phi(x) = \left[\begin{array}{ll} x \log x - x + 1 & \text{if } x \leq e, \\ \frac{1}{2e}(x-e)^2 + (x-e) + 1 & \text{if } x > e. \end{array} \right.$$

We use this divergence because Assumption Φ(ii) fails for KL divergence, whereas hybrid divergence only requires finite second moments for Assumption Φ(ii). The LFDs under hybrid divergence are also everywhere positive, which is not guaranteed under χ^2 divergence. We repeated our analysis with neighborhoods constrained by χ^2 and L^4 divergences as robustness checks. Overall, our findings are not sensitive to ϕ (see Appendix D.1 for a discussion).

Findings. Figure 1 presents a sensitivity analysis of the "some college" to "college graduate" premium. Cohort-wise estimates and CSs for $\underline{\kappa}_\delta$ and $\overline{\kappa}_\delta$ are presented, beginning at $\delta = 0.01$ and increasing δ by factors of 10 up to $\delta = 100$. Even with $\delta = 0.01$, estimates of $\underline{\kappa}_\delta$ and $\overline{\kappa}_\delta$ lie uniformly below and above zero across cohorts without exchangeability (see Figure 1a). Imposing exchangeability can tighten the bounds, with the bounds for $\delta = 0.01$ significantly negative in early cohorts and significantly positive in later cohorts (see Figure 1b). But the $\delta = 0.1$ bounds with exchangeability again contain zero across all cohorts. Bounds for larger δ presented in Figures 1c and 1d are uninformatively wide.

To understand better what is meant by "small" and "large" neighborhoods, Figure 2 plots marginal CDFs for the LFDs under which the upper bounds for cohort 1 are attained. Similar LFDs (not reported) were obtained for other cohorts and the lower bounds. Without exchangeability, the LFDs with $\delta = 0.1$ are almost identical to Gumbel
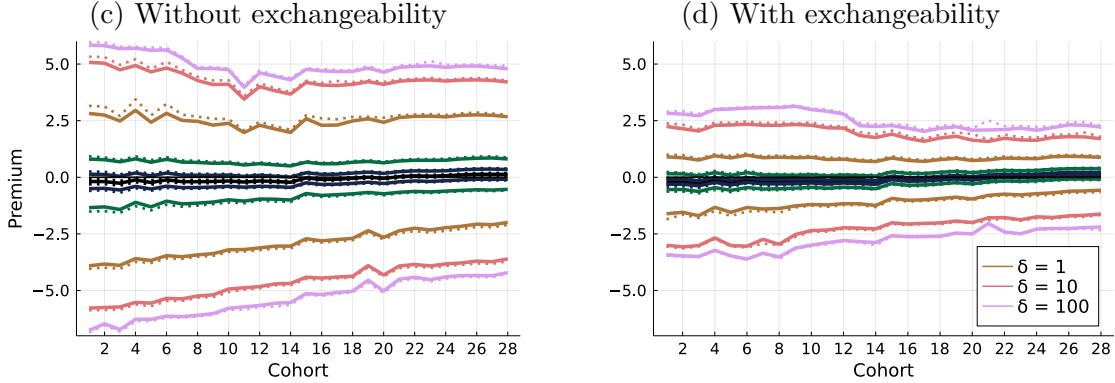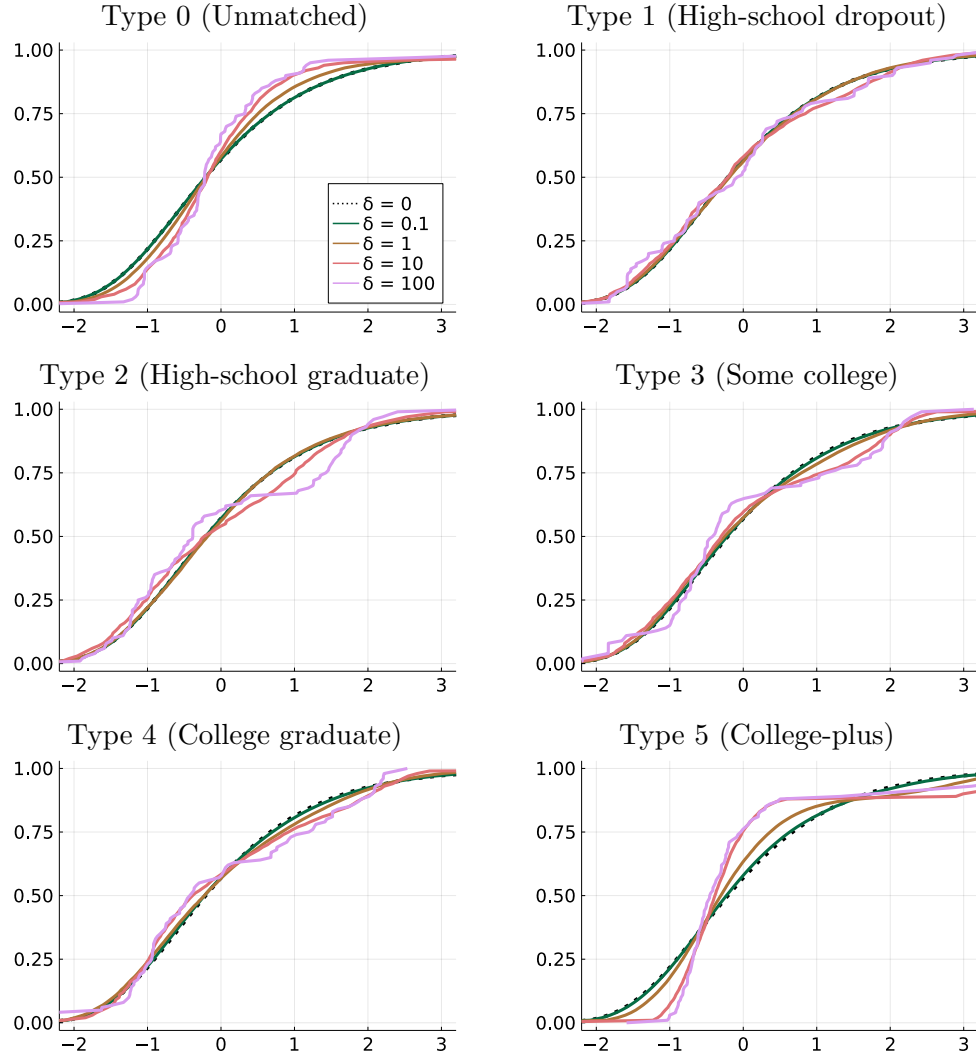
Smaller neighborhoods
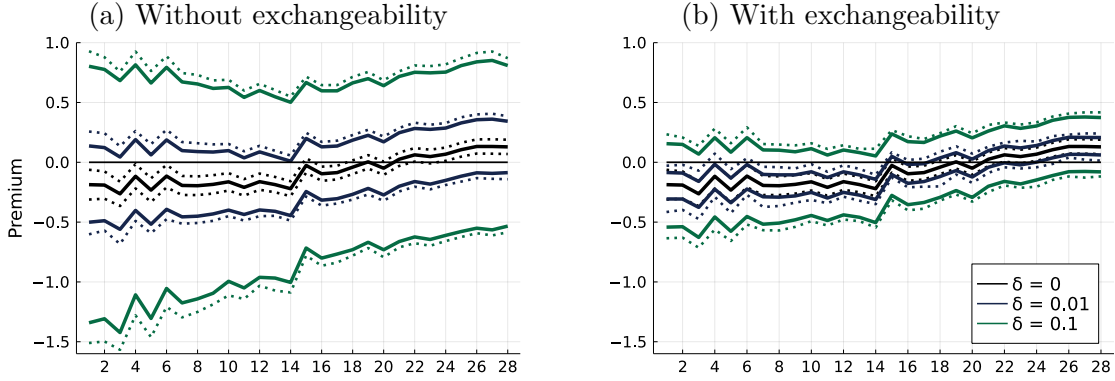
Larger neighborhoods

Figure 1: Sensitivity analysis of the "some college" to "college graduate" premium across cohorts. *Note:* Solid lines are estimates, dotted lines are (cohort-wise) 95% CSs. CSW's estimates and CSs correspond to $\delta = 0$.

(plots with $\delta = 0.01$ are indistinguishable from Gumbel). LFDs appear close to Gumbel across most potential spouse types with $\delta = 1$, while for $\delta = 10$ and $\delta = 100$ the LFDs have kinks and indicate shifts in mass from the center of the distribution to the tails.

Under exchangeability (Figure 2b), the marginal distribution of shocks is independent of potential spouse type. In this case the LFDs for $\delta = 1$ or smaller are virtually indistinguishable from Gumbel. LFDs with $\delta = 10$ and $\delta = 100$ are also less kinked than Figure 2a because distortions are spread more evenly across potential spouse types.

We also computed the largest correlation of shocks under the LFDs at which the bounds are attained and our *size* measure from Section 4.3. As these quantities are stable across cohorts, we present their averages across cohorts in Table 3. Shocks are

29

Type 0 (Unmatched)

Type 1 (High-school dropout)

Type 2 (High-school graduate)

Type 3 (Some college)

Type 4 (College graduate)

Type 5 (College-plus)

(b) With exchangeability (all types)

Figure 2: Marginal CDFs for the LFDs maximizing the "some college" to "college grad-uate" premium in cohort 1 across potential spouse types.

independent when $\delta = 0$ and only very weakly correlated for small δ, while for large δ some shocks are strongly negatively correlated. The maximal correlations under exchangeability are smaller, especially for large δ. Turning to the *size* measure, the LFDs for $\delta = 0.01$ without exchangeability shift the model-implied marriage probabilities by 0.01 (on average, across cohorts) from their values under the i.i.d. Gumbel assumption. LFDs for $\delta = 10$ and $\delta = 100$ shift marriage probabilities around 0.25 (on average, across cohorts). Imposing exchangeability reduces the *size* measure by around 25% because model parameters do not vary as much under this shape restriction.

In view of the small-δ bounds in Figure 1, the LFDs in Figure 2, and the metrics in Table 3, it seems impossible to draw conclusions about how the sign of the premium has changed over time under slight nonparametric relaxations of the i.i.d. Gumbel assumption. To help understand why, Figure 5 plots bounds where F is allowed to vary but θ is held fixed at CSW's estimates. These "fixed-θ" bounds for $\delta = 10$ and $\delta = 100$ are almost identical, and are roughly the same width as the $\delta = 0.01$ bounds in Figure 1. The width of the bounds in Figure 1 therefore seems largely attributable to the additional variation in θ that is permitted when parametric assumptions for F are relaxed.

Overall, our findings are complementary to Gualdani and Sinha (2020) who perform a nonparametric reanalysis of CSW using the PIES methodology of Torgovitsky (2019b). Although they do not derive nonparametric bounds on the marital education premium itself, only terms that contribute to it, they also find no evidence of an increase in premiums across cohorts.

5.2 Welfare Analysis in a Rust Model

Our second empirical illustration is a sensitivity analysis for welfare counterfactuals in the DDC model of Rust (1987).

Model and Benchmark Estimates. We focus on the specification in Table IX of Rust (1987) where maintenance costs are linear in the state (i.e., mileage). In the notation of Example 2.3, $|\mathcal{S}| = 90$, $\beta = 0.9999$, and $\theta_\pi = (RC, MC)$ where RC is the replacement cost and MC is a maintenance cost parameter. Our counterfactual of interest is the change in average welfare arising from a 10% reduction in maintenance costs. Hence, $\pi_{1,s}(\theta_\pi) = \tilde{\pi}_{1,s}(\theta_\pi) = -RC$ and $\pi_{0,s}(\theta_\pi) = -0.001MC \times s$ (baseline) and $\tilde{\pi}_{0,s}(\theta_\pi) = 0.9\pi_{0,s}(\theta_\pi)$ (counterfactual). The counterfactual function is $k(\theta, \gamma) = w'(\tilde{v} - v)$ where w is the stationary distribution of the state in the baseline model.

Under the i.i.d. Gumbel assumption, the estimated counterfactual at the maximum

likelihood estimate (MLE) of θ_π is 73.07 and its 95% CS is [48.25,101.31].[21] Note the counterfactual is point-identified under the i.i.d. Gumbel assumption because θ_π is point-identified.

Implementation. We estimate CCPs using Rust's Group 4 data. Nonparametric estimates of the 90 CCPs are zero in many states, so we proceed as in Section 3.3 and take the model-implied CCPs at the MLE of θ_π (under the i.i.d. Gumbel assumption) as our estimate $\hat{P}_2$. We drop moment conditions for CCPs in states where the replacement probability is less than 0.001 to avoid numerical instabilities induced by including near-degenerate moments. This reduces the dimension of g_2 to 71. We normalize F so that shocks have mean zero and the same variance as the Gumbel distribution by appending $\mathbb{E}^F[U_d] = 0$ and $\mathbb{E}^F[U_d^2 - \pi^2/6] = 0$, for $d = 0, 1$, to g_4. In total, there are 255 moments (71 for CCPs, 180 for Bellman equations, and 4 location/scale normalizations) and $\theta = (\theta_\pi, v, \tilde{v})$ has dimension 182.

We implement our methods as described in Section 3.2. The inner optimization uses 75 moments (71 for CCPs and 4 for normalizations), with the remaining 180 moments appended as constraints in the outer optimization. We define neighborhoods using hybrid divergence from Section 5.1 so that Assumption Φ(ii) holds. Similar results are obtained with χ^2 and L^4 neighborhoods (see Appendix D.2). Expectations are computed using 50,000 scrambled Halton draws—see Appendix D.2 for computation times. We compute 95% CSs for $\underline{\kappa}_\delta$ and $\overline{\kappa}_\delta$ using the bootstrap procedure from Section 6.2 and projection procedure from Section 6.3. See Appendix D.2 for details.

Findings. Estimates and CSs for $\underline{\kappa}_\delta$ and $\overline{\kappa}_\delta$ are plotted in Figure 3 for values of δ from 0.01 to 100.[22] As can be seen, the bounds expand rapidly under slight relaxations of the i.i.d. Gumbel assumption then stabilize around $\delta = 1$, where the lower bound is 6.45 and the upper bound of 160.5 represents approximately 220% of the value under the i.i.d. Gumbel assumption.

To interpret δ, in Figure 4 we plot the CDFs of $U_1 - U_0$ under the LFDs at which the estimated bounds $\underline{\hat{\kappa}}_\delta$ and $\overline{\hat{\kappa}}_\delta$ are attained. LFDs were computed as described in

---

[21]We construct this CS by simulation. We draw $\hat{\theta}_\pi^* \sim N(\hat{\theta}_\pi, \hat{\Sigma})$ where $\hat{\theta}_\pi$ is the MLE and $\hat{\Sigma}$ is an estimate of the inverse information matrix. For each $\hat{\theta}_\pi^*$ draw, we compute the baseline and counterfactual value functions v^* and $\tilde{v}^*$, and hence the counterfactual $\hat{\kappa}^* = w'(\tilde{v}^* - v^*)$.

[22]The width of the bootstrap CSs relative to the bounds reduces as δ gets large. We re-estimated our bounds using several different draws of bootstrapped CCPs in place of $\hat{P}_2$ and obtained bounds that spanned a range similar to the bootstrap CSs for small δ, but which for many draws converged to values close to our estimates of the bounds for large δ. This corroborates the behavior of our bootstrap CSs. We conjecture that other features of the model are potentially more important than the numerical values of the CCPs in determining nonparametric bounds on the welfare counterfactual.

Figure 3: Sensitivity analysis for change in average welfare under a 10% maintenance cost subsidy. *Note:* Solid lines are estimates, dotted lines are bootstrap CSs, dashed lines are projection CSs.

Section 4.2 using the construction (27). The distributions appear very close to logistic (their distribution when $\delta = 0$) for $\delta = 0.01$. Therefore, we see that large differences in welfare counterfactuals can arise under very slight departures from the i.i.d. Gumbel assumption. LFDs for the upper bound shift increasing amounts of mass to the center of the distribution of $U_1 - U_0$ as δ increases. LFDs corresponding to the lower bound are relatively less distorted, but have increasing amounts of mass shifted into the right tail. These are similar for $\delta = 0.1$ through $\delta = 100$ because the estimated lower bound stabilizes for smaller values of δ than the upper bound (cf. Figure 3).

Table 4 lists other metrics to help interpret the neighborhood size. The first is the correlation of U_0 and U_1 under the LFDs at which $\hat{\underline{\kappa}}_\delta$ and $\hat{\overline{\kappa}}_\delta$ are attained. These are very small for $\delta = 0.01$ and remain small under the LFDs for $\hat{\underline{\kappa}}_\delta$ as δ increases, while U_0 and U_1 are strongly positively correlated under the LFDs for $\hat{\overline{\kappa}}_\delta$, especially for larger δ values. Given the asymmetry in distortions between the lower and upper values, we compute our *size* measure separately for both. We measure distortions to the moments corresponding to the CCPs as these are most directly interpretable within the context of the model. We see that the LFDs for $\delta = 0.01$ are distorting F_* in a manner that shifts the model-implied CCPs by at most 0.016. By contrast, the LFDs for $\delta = 10$ and $\delta = 100$ shift the model-implied CCPs from their values under the i.i.d. Gumbel assumption by at most 0.04 for $\hat{\underline{\kappa}}_\delta$ and 0.47 for $\hat{\overline{\kappa}}_\delta$.

The parameters at which $\hat{\underline{\kappa}}_\delta$ and $\hat{\overline{\kappa}}_\delta$ are attained are also revealing about neighborhood size. Table 4 presents MLEs of MC and RC, which are similar to the values

33

Figure 4: CDFs of $U_1 - U_0$ under the LFDs at which the estimated lower and upper bounds on the welfare counterfactual are attained.

Table 4: Metrics for interpreting δ

	Lower bound				Upper bound			
δ	$corr$	$size$	RC	MC	$corr$	$size$	RC	MC
0	0.000	0.000	10.208	2.294	0.000	0.000	10.208	2.294
0.01	0.036	0.010	7.357	1.411	-0.027	0.016	13.390	3.307
0.1	-0.058	0.039	5.186	0.553	0.149	0.109	16.134	4.374
1	-0.045	0.039	4.023	0.203	0.616	0.346	17.166	5.038
10	-0.040	0.039	4.022	0.202	0.765	0.461	17.595	5.331
100	-0.063	0.039	3.931	0.176	0.764	0.469	17.626	5.365

Note: Correlation of U_0 and U_1 under the LFD at which the estimated lower and upper bounds are attained (*corr*), our *size* measure from Section 4.3, and replacement and maintenance cost parameters at which the estimated lower and upper bounds are attained.

reported in Table IX of Rust (1987). We see from Table 4 that $\hat{\underline{\kappa}}_\delta$ and $\hat{\bar{\kappa}}_\delta$ are attained at very different parameter values, with much smaller cost parameters for the lower bound and larger parameters for the upper bound, even for $\delta = 0.01$. Intuitively, a smaller MC means that the saving from the subsidy—which is proportional—must be small. Correspondingly, a low RC is needed to help the model to fit the observed CCPs at the smaller MC. While it is known that payoff parameters are not identified without parametric assumptions on F, it is perhaps surprising that these parameters vary by so much under slight relaxations of the i.i.d. Gumbel assumption. For instance, with $\delta = 0.01$ the lower bound is attained with cost parameters $RC = 7.357$ and $MC = 1.411$ while the upper bound is attained with cost parameters that are roughly double these values.

6 Estimation and Inference

We begin in Section 6.1 by establishing consistency and the asymptotic distribution of the estimators $\hat{\underline{\kappa}}_\delta$ and $\hat{\bar{\kappa}}_\delta$ from Section 2.4. We then present a bootstrap-based inference method in Section 6.2 and a projection-based inference method in Section 6.3.

6.1 Large-sample Properties of Plug-in Estimators

We first introduce some regularity conditions. Recall the space $\mathcal{E}$ from Assumption Φ. We equip $\mathcal{E}$ with the Orlicz norm (see Appendix F of Christensen and Connault (2022))

$$\|f\|_\psi = \inf_{c>0} \frac{1}{c} \left(1 + \mathbb{E}^{F_*}[\psi(c|f(U)|)]\right).$$

This norm is equivalent to the $L^2(F_*)$ norm for χ^2 and hybrid divergence and equivalent to the $L^q(F_*)$ norm for L^p divergence ($p^{-1} + q^{-1} = 1$), while for KL divergence it is stronger than any $L^p(F_*)$ norm with $p < \infty$ but weaker than the sup-norm. Say that a class of functions $\{f_a : a \in \mathcal{A}\} \subset \mathcal{E}$ indexed by a metric space $\mathcal{A}$ is $\mathcal{E}$-*continuous in* a if $a' \to a$ in $\mathcal{A}$ implies $\|f_a - f_{a'}\|_\psi \to 0$. We also require a slightly stronger notion of constraint qualification than Condition S from Section 2.5.

Definition 6.1 Condition S' *holds at* (θ, γ, P) *if* $\vec{P} \in \text{int}(\mathcal{G}(\theta, \gamma) + \mathcal{C})$.

Condition S' replaces "relative interior" in Condition S with "interior". Finally, recall $\Delta(\theta; \gamma, P)$ from (16) and let $\Theta_\delta(\gamma, P) = \{\theta \in \Theta : \Delta(\theta; \gamma, P) < \delta\}$.

Assumption M *(i)* $k(\cdot; \theta, \gamma)$ *and each entry of* $g(\cdot; \theta, \gamma)$ *are* $\mathcal{E}$-*continuous in* (θ, γ);
 (ii) $(\theta, \gamma) \mapsto \mathbb{E}^{F_*}[\phi^\star(a_1 + a_2 k(U, \theta, \gamma) + a_3' g(U, \theta, \gamma))]$ *is continuous for each* $(a_1, a_2, a_3) \in \mathbb{R} \times \mathbb{R} \times \mathbb{R}^d$;
 (iii) $\Theta_\delta(\gamma_0, P_0)$ *is nonempty and Condition S' holds at* (θ, γ_0, P_0) *for each* $\theta \in \Theta_\delta(\gamma_0, P_0)$;
 (iv) $\text{cl}(\Theta_\delta(\gamma_0, P_0)) \supseteq \{\theta \in \Theta : \Delta(\theta; \gamma_0, P_0) \leq \delta\}$;
 (v) Θ *is a compact subset of* $\mathbb{R}^{d_\theta}$.

Parts (i) and (ii) of Assumption M are continuity conditions. If k and g consist of indicator functions, then these conditions hold provided the probabilities of the events under F_* are continuous in (θ, γ). In models without γ, these conditions simply require continuity in θ.

There are two parts to Assumption M(iii). The nonemptyness condition holds when the model is correctly specified under F_* or, more generally, when there is at least one

$F \in \mathcal{N}_\delta$ that satisfies (1) for some θ. The second part is a constraint qualification. This condition requires that for each $\theta \in \Theta_\delta(\gamma_0, P_0)$, there is a distribution F under which (1) holds at (θ, γ_0, P_0) that is "interior" to $\mathcal{N}_\infty$, in the sense that one can perturb the moments at (θ, γ_0, P_0) in all directions by perturbing F. Condition S' also requires that there is $F \in \mathcal{N}_\infty$ under which any inequality restrictions at (θ, γ_0, P_0) hold strictly. Note, however, that we do not require that this F belongs to $\mathcal{N}_\delta$, only to $\mathcal{N}_\infty$. We therefore do not view this condition as overly restrictive. We also conjecture it could be relaxed using a notion similar to S-regularity from Section 2.5.

Assumption M(iv) is made for convenience and can be relaxed; this condition simply ensures that there do not exist values of θ at which $\Delta(\theta; \gamma_0, P_0) = \delta$ that are separated from $\Theta_\delta(\gamma_0, P_0)$. Assumption M(v) is standard and can be relaxed.

Theorem 6.1 *Suppose that Assumptions Φ and M hold and $(\hat\gamma, \hat P) \to_p (\gamma_0, P_0)$ or, if there is no auxiliary parameter, $\hat P \to_p P_0$. Then $\hat{\underline\kappa}_\delta \to_p \underline\kappa_\delta$ and $\hat{\overline\kappa}_\delta \to_p \overline\kappa_\delta$.*

To derive the asymptotic distribution of the estimators, we assume γ_0 is known and suppress dependence of all quantities on γ for the remainder of this section. This entails no loss of generality for models without γ, such as Examples 2.1 and 2.2 and the application in Section 5.1. In DDC models this presumes the law of motion of the state is known. The asymptotic distribution therefore reflects only sampling uncertainty from the estimated CCPs, which is the case for confidence sets reported when laws of motion are first estimated "offline". Extending our approach to accommodate sampling variation in $\hat\gamma$ in a tractable manner appears to require exploiting application-specific model structure, which we defer to future work.

Define

$$\underline{b}_\delta(P) = \inf_{\theta \in \Theta_\delta(P)} \underline{K}_\delta(\theta; P), \qquad \overline{b}_\delta(P) = \sup_{\theta \in \Theta_\delta(P)} \overline{K}_\delta(\theta; P). \tag{29}$$

In this notation, $\underline\kappa_\delta = \underline{b}_\delta(P_0)$ and $\overline\kappa_\delta = \overline{b}_\delta(P_0)$ (see Lemma E.3) and $\hat{\underline\kappa}_\delta = \underline{b}_\delta(\hat P)$ and $\hat{\overline\kappa}_\delta = \overline{b}_\delta(\hat P)$. We derive the asymptotic distribution of $\hat{\underline\kappa}_\delta$ and $\hat{\overline\kappa}_\delta$ by showing $\underline{b}_\delta$ and $\overline{b}_\delta$ are directionally differentiable and applying a suitable delta method. Say $f : \mathbb{R}^{d_1+d_2} \to \mathbb{R}$ is (Hadamard) directionally differentiable at P_0 if there is a continuous map $df_{P_0}[\cdot] : \mathbb{R}^{d_1+d_2} \to \mathbb{R}$ such that

$$\lim_{n \to \infty} t_n^{-1} \left(f(P_0 + t_n h_n) - f(P_0) \right) = df_{P_0}[h]$$

for all sequences $t_n \downarrow 0$ and $h_n \to h$ (Shapiro, 1990, p. 480). If $df_{P_0}[h]$ is linear in h then f is (fully) differentiable at P_0. We introduce some additional notation used to define

the directional derivatives of $\underline{b}_\delta$ and $\overline{b}_\delta$. Let

$$\underline{\Xi}_\delta(\theta; P) = \mathrm{argsup}_{\eta \geq 0, \zeta \in \mathbb{R}, \lambda \in \Lambda} - \mathbb{E}^{F_*}\left[(\eta\phi)^\star(-k(U, \theta) - \zeta - \lambda' g(U, \theta))\right] - \eta\delta - \zeta - \lambda'_{12} P,$$

where $(\eta\phi)^\star$ denotes the convex conjugate of $x \mapsto \eta \cdot \phi(x)$, and let $\overline{\Xi}_\delta(\theta; P)$ denote the analogous arginf for the minimization problem corresponding to the upper bound. Recall that $\underline{\lambda}_{12} = (\underline{\lambda}_1, \underline{\lambda}_2)$ collects the first $d_1 + d_2$ elements of $\underline{\lambda}$. Let

$$\underline{\Lambda}_\delta(\theta; P) = \{(\lambda_1, \lambda_2) : (\eta, \zeta, \lambda_1, \lambda_2, \lambda_3, \lambda_4) \in \underline{\Xi}_\delta(\theta; P)\}$$

denote the projection of $\underline{\Xi}_\delta(\theta; P)$ for $\underline{\lambda}_{12}$. We let $\overline{\Lambda}_\delta(\theta; P)$ denoting the analogous projection of $\overline{\Xi}_\delta(\theta; P)$. Finally, let

$$\underline{\Theta}_\delta(P_0) = \arg\min_{\theta \in \Theta} \underline{K}_\delta(\theta; P_0), \quad \overline{\Theta}_\delta(P_0) = \arg\max_{\theta \in \Theta} \overline{K}_\delta(\theta; P_0).$$

The sets $\underline{\Theta}_\delta(P_0)$ and $\overline{\Theta}_\delta(P_0)$ are nonempty and compact under Assumptions Φ and M.

The following regularity conditions are presented for the general case where k depends on u. It may be possible to weaken some of these regularity conditions in the special case in which k does not depend on u.

Assumption M (continued) *(vi)* $\underline{\Theta}_\delta(P_0) \subseteq \Theta_\delta(P_0)$ *and* $\overline{\Theta}_\delta(P_0) \subseteq \Theta_\delta(P_0)$;
(vii) $\theta \mapsto \underline{\Lambda}_\delta(\theta; P_0)$ *and* $\theta \mapsto \overline{\Lambda}_\delta(\theta; P_0)$ *are lower hemicontinuous at each* $\theta \in \underline{\Theta}_\delta(P_0)$ *and* $\theta \in \overline{\Theta}_\delta(P_0)$, *respectively.*

Theorem 6.2 *Suppose that Assumptions Φ and M hold. Then $\underline{b}_\delta$ and $\overline{b}_\delta$ are directionally differentiable at P_0, with*

$$d\underline{b}_{\delta, P_0}[h] = \min_{\theta \in \underline{\Theta}_\delta(P_0)} \max_{\underline{\lambda}_{12} \in \underline{\Lambda}_\delta(\theta; P_0)} -\underline{\lambda}'_{12} h, \qquad d\overline{b}_{\delta, P_0}[h] = \max_{\theta \in \overline{\Theta}_\delta(P_0)} \min_{\overline{\lambda}_{12} \in \overline{\Lambda}_\delta(\theta; P_0)} \overline{\lambda}'_{12} h.$$

Moreover, if $\sqrt{n}(\hat{P} - P_0) \to_d Z \sim N(0, \Sigma)$ with Σ finite, then

$$\sqrt{n}\left(\begin{pmatrix} \hat{\overline{\kappa}}_\delta \\ \hat{\underline{\kappa}}_\delta \end{pmatrix} - \begin{pmatrix} \overline{\kappa}_\delta \\ \underline{\kappa}_\delta \end{pmatrix}\right) \to_d \begin{pmatrix} d\underline{b}_{\delta, P_0}[Z] \\ d\overline{b}_{\delta, P_0}[Z] \end{pmatrix}.$$

The asymptotic distribution presented in Theorem 6.2 is non-Gaussian. In the special case in which $\cup_{\theta \in \underline{\Theta}_\delta(P_0)} \underline{\Lambda}_\delta(\theta; P_0) = \{\underline{\lambda}_{12}\}$, the asymptotic distribution of $\hat{\underline{\kappa}}_\delta$ simplifies to $N(0, \underline{\lambda}'_{12} \Sigma \underline{\lambda}_{12})$. An analogous simplification holds for $\hat{\overline{\kappa}}_\delta$ when $\cup_{\theta \in \overline{\Theta}_\delta(P_0)} \overline{\Lambda}_\delta(\theta; P_0)$ is a singleton.

6.2 Inference Procedure 1: Bootstrap

Our first inference procedure specializes the general approach of Fang and Santos (2019) for inference on directionally differentiable functions to the present setting. Define

$$\widehat{d\underline{b}}_{\delta,P_0}[h] = \inf_{\theta \in \hat{\underline{\Theta}}_{\delta,n}} \sup_{\underline{\lambda}_{12} \in \underline{\Lambda}_\delta(\theta;\hat{P})} -\underline{\lambda}'_{12}h, \qquad \widehat{d\overline{b}}_{\delta,P_0}[h] = \sup_{\theta \in \hat{\overline{\Theta}}_{\delta,n}} \inf_{\overline{\lambda}_{12} \in \overline{\Lambda}_\delta(\theta;\hat{P})} \overline{\lambda}'_{12}h,$$

where

$$\hat{\underline{\Theta}}_{\delta,n} = \{\theta \in \Theta_\delta(\hat{P}) : \underline{K}_\delta(\theta;\hat{P}) \leq \hat{\underline{\kappa}}_\delta + \hat{\nu}\sqrt{\log n/n}\}, \text{ and}$$

$$\hat{\overline{\Theta}}_{\delta,n} = \{\theta \in \Theta_\delta(\hat{P}) : \overline{K}_\delta(\theta;\hat{P}) \geq \hat{\overline{\kappa}}_\delta - \hat{\nu}\sqrt{\log n/n}\},$$

with $\hat{\nu}$ a (possibly random) positive scalar tuning parameter for which $\hat{\nu} \to_p \nu > 0$. Any such $\hat{\nu}$ results in a confidence set with asymptotically correct coverage. We give some practical guidance for choosing $\hat{\nu}$ below.

Let $\hat{P}^*$ denote a bootstrapped version of $\hat{P}$. In practice any bootstrap can be used provided it satisfies mild consistency conditions. In the empirical application in Section 5.1 we simply draw $\hat{P}^* \sim N(\hat{P}, \hat{\Sigma}/n)$ where $\hat{\Sigma}$ is a consistent estimator of Σ. Let

$$\hat{\underline{c}}_\alpha = \alpha\text{-quantile of } \widehat{d\underline{b}}_{\delta,P_0}[\sqrt{n}(\hat{P}^* - \hat{P})], \quad \hat{\overline{c}}_\alpha = \alpha\text{-quantile of } \widehat{d\overline{b}}_{\delta,P_0}[\sqrt{n}(\hat{P}^* - \hat{P})],$$

where the quantiles are computed by resampling $\hat{P}^*$ (conditional on the data). Lower, upper, and two-sided $100(1 - \alpha)\%$ CSs for $\underline{\kappa}_\delta$ and $\overline{\kappa}_\delta$ are

$$CS_{\delta,L}^{1-\alpha} = \left[\hat{\underline{\kappa}}_\delta - \frac{\hat{\underline{c}}_{1-\alpha}}{\sqrt{n}}, +\infty\right),$$

$$CS_{\delta,U}^{1-\alpha} = \left(-\infty, \hat{\overline{\kappa}}_\delta - \frac{\hat{\overline{c}}_\alpha}{\sqrt{n}}\right], \qquad CS_\delta^{1-\alpha} = \left[\hat{\underline{\kappa}}_\delta - \frac{\hat{\underline{c}}_{1-\alpha/2}}{\sqrt{n}}, \hat{\overline{\kappa}}_\delta - \frac{\hat{\overline{c}}_{\alpha/2}}{\sqrt{n}}\right].$$

We require a slight strengthening of Assumption M(vii) to establish validity of the procedure. As before, regularity conditions are presented for the general case where k depends on u. It may be possible to weaken these conditions when k does not depend on u.

Assumption M (continued) *(vii') $(\theta, P) \mapsto \underline{\Lambda}_\delta(\theta; P)$ and $(\theta, P) \mapsto \overline{\Lambda}_\delta(\theta; P)$ are lower hemicontinuous at (θ, P_0) for each $\theta \in \underline{\Theta}_\delta(P_0)$ and $\theta \in \overline{\Theta}_\delta(P_0)$, respectively.*

Theorem 6.3 *Suppose that Assumptions Φ and M(i)–(vi),(vii') hold, $\sqrt{n}(\hat{P} - P_0) \to_d Z \sim N(0, \Sigma)$ with Σ finite, and $\hat{P}^*$ satisfies Assumption 3 of Fang and Santos (2019).*

Then the distribution of $\widehat{d\underline{b}}_{\delta,P_0}[\sqrt{n}(\hat{P}^* - \hat{P})]$ and $\widehat{d\overline{b}}_{\delta,P_0}[\sqrt{n}(\hat{P}^* - \hat{P})]$ (conditional on the data) is consistent for the asymptotic distribution derived in Theorem 6.2. Moreover, if the CDFs of $d\underline{b}_{\delta,P_0}[Z]$ and $d\overline{b}_{\delta,P_0}[Z]$ are continuous and increasing at their $\alpha/2$, α, $1-\alpha$, and $1 - \alpha/2$ quantiles, then

$$\lim_{n\to\infty} \Pr(\underline{\kappa}_\delta \in CS_{\delta,L}^{1-\alpha}) = 1 - \alpha\,,$$

$$\lim_{n\to\infty} \Pr(\overline{\kappa}_\delta \in CS_{\delta,U}^{1-\alpha}) = 1 - \alpha\,, \qquad \liminf_{n\to\infty} \Pr([\underline{\kappa}_\delta, \overline{\kappa}_\delta] \subseteq CS_\delta^{1-\alpha}) \geq 1 - \alpha\,.$$

Any $\hat{\nu}$ that satisfies $\hat{\nu} \to_p \nu > 0$ results in asymptotically valid CSs. In view of the functional forms of $\widehat{d\underline{b}}_{\delta,P_0}[\cdot]$ and $\widehat{d\overline{b}}_{\delta,P_0}[\cdot]$, smaller $\hat{\nu}$ produce (weakly) wider CSs. In the CSW application, we set $\hat{\nu}$ equal to the minimum diagonal element of the covariance matrix of the moments evaluated at $(\hat{\theta}, \hat{\gamma}, \hat{P})$ under F_*, where $\hat{\theta}$ is computed under F_*. We chose this quantity as it is related to the convexity of the inner problem for small δ. In practice, this resulted in $\hat{\nu}$ between 0.001 and 0.01. We recommend setting $\hat{\nu}$ to be of a similarly small magnitude, then performing a sensitivity analysis to check that critical values aren't too dependent on $\hat{\nu}$. Setting $\hat{\nu} = 0$ and replacing $\widehat{\underline{\Theta}}_{\delta,n}$ and $\widehat{\overline{\Theta}}_{\delta,n}$ by $\{\hat{\underline{\theta}}_\delta\}$ and $\{\hat{\overline{\theta}}_\delta\}$ where $\hat{\underline{\theta}}_\delta$ and $\hat{\overline{\theta}}_\delta$ minimize and maximize the sample criterions is also valid, but may be conservative.

6.3 Inference Procedure 2: Projection

This second approach is computationally simple but possibly conservative.[23] Suppose we have random vectors $\hat{P}_{1,U}^{1-\alpha}$, $\hat{P}_{2,U}^{1-\alpha}$, and $\hat{P}_{2,L}^{1-\alpha}$ that form a $100(1-\alpha)\%$ rectangular CS for P_0:

$$\liminf_{n\to\infty} \Pr\left(P_{10} \leq \hat{P}_{1,U}^{1-\alpha}\,, \hat{P}_{2,L}^{1-\alpha} \leq P_{20} \leq \hat{P}_{2,U}^{1-\alpha}\right) \geq 1 - \alpha\,, \tag{30}$$

where the inequalities should be understood to hold element-wise (we discuss how to construct a rectangular CS for P_0 below).

The idea behind this approach is to replace any moment conditions involving P by inequalities constructed from the rectangular CS. Define the criterion functions

$$\hat{\underline{K}}_{\delta,1-\alpha}(\theta) = \begin{bmatrix} \underline{K}_{\delta,cs}(\theta; \hat{P}_{1-\alpha}) \\ +\infty \end{bmatrix}\,, \quad \hat{\overline{K}}_{\delta,1-\alpha}(\theta) = \begin{bmatrix} \overline{K}_{\delta,cs}(\theta; \hat{P}_{1-\alpha}) & \text{if } \Delta_{cs}(\theta; \hat{P}_{1-\alpha}) < \delta, \\ -\infty & \text{if } \Delta_{cs}(\theta; \hat{P}_{1-\alpha}) \geq \delta, \end{bmatrix}$$

---

where $\underline{K}_{\delta,cs}$, $\overline{K}_{\delta,cs}$, and Δ_{cs} are versions of (13), (14), and (16) formed using

$$\mathbb{E}^F[g_1(U,\theta)] \leq \hat{P}_{1,U}^{1-\alpha}\,,\quad \mathbb{E}^F[g_2(U,\theta)] \leq \hat{P}_{2,U}^{1-\alpha}\,,\quad \mathbb{E}^F[-g_2(U,\theta)] \leq -\hat{P}_{2,L}^{1-\alpha}\,,\tag{31}$$

as well as (1c) and (1d). In these criterions, Λ is replaced by $\Lambda_{cs} = \mathbb{R}_+^{d_1+2d_2+d_3} \times \mathbb{R}^{d_4}$, g is replaced by $g_{cs} = (g_1, g_2, -g_2, g_3, g_4)$, P is replaced by $\hat{P}_{1-\alpha} = (\hat{P}_{1,U}^{1-\alpha}, \hat{P}_{2,U}^{1-\alpha}, -\hat{P}_{2,L}^{1-\alpha})$, and λ_{12} denotes the first $d_1 + 2d_2$ elements of λ.

Critical values are computed by optimizing the criterions $\hat{\underline{K}}_{\delta,1-\alpha}$ and $\hat{\overline{K}}_{\delta,1-\alpha}$ with respect to θ:

$$\hat{\underline{\kappa}}_{\delta,1-\alpha} = \inf_{\theta\in\Theta} \hat{\underline{K}}_{\delta,1-\alpha}(\theta)\,,\quad \hat{\overline{\kappa}}_{\delta,1-\alpha} = \sup_{\theta\in\Theta} \hat{\overline{K}}_{\delta,1-\alpha}(\theta)\,.$$

Lower, upper, and two-sided $100(1-\alpha)\%$ CSs for $\underline{\kappa}_\delta$ and $\overline{\kappa}_\delta$ are then given by

$$CS_{\delta,L}^{1-\alpha} = \left[\hat{\underline{\kappa}}_{\delta,1-\alpha}, +\infty\right)\,,\quad CS_{\delta,U}^{1-\alpha} = \left(-\infty, \hat{\overline{\kappa}}_{\delta,1-\alpha}\right]\,,\quad CS_\delta^{1-\alpha} = \left[\hat{\underline{\kappa}}_{\delta,1-\alpha}, \hat{\overline{\kappa}}_{\delta,1-\alpha}\right]\,.$$

Theorem 6.4 *Suppose that Assumptions Φ and M(i),(iii)–(v) hold and $\hat{P}_{1-\alpha}$ satisfies (30). Then*

$$\liminf_{n\to\infty} \Pr(\underline{\kappa}_\delta \in CS_{\delta,L}^{1-\alpha}) \geq 1-\alpha\,,$$

$$\liminf_{n\to\infty} \Pr(\overline{\kappa}_\delta \in CS_{\delta,U}^{1-\alpha}) \geq 1-\alpha\,,\qquad \liminf_{n\to\infty} \Pr([\underline{\kappa}_\delta, \overline{\kappa}_\delta] \subseteq CS_\delta^{1-\alpha}) \geq 1-\alpha\,.$$

To construct a rectangular CS for P_0 satisfying (30), suppose $\sqrt{n}(\hat{P}-P_0) \to_d N(0,\Sigma)$ and we have a consistent estimator $\hat{\Sigma}$ of Σ. Let $\hat{\sigma}$ denote the vector formed by taking the square root of each diagonal entry of $\hat{\Sigma}$. Partition $\hat{\sigma}$ conformably as $\hat{\sigma} = (\hat{\sigma}_{(1)}, \hat{\sigma}_{(2)})$ and set

$$\hat{P}_{1,L}^{1-\alpha} = \hat{P}_1 + n^{-1/2}\hat{c}_{1-\alpha,1}\hat{\sigma}_{(1)}\,,\quad \hat{P}_{2,L}^{1-\alpha} = \hat{P}_2 - n^{-1/2}\hat{c}_{1-\alpha,2}\hat{\sigma}_{(2)}\,,\quad \hat{P}_{2,U}^{1-\alpha} = \hat{P}_2 + n^{-1/2}\hat{c}_{1-\alpha,2}\hat{\sigma}_{(2)}\,,$$

where the (scalar) critical values $\hat{c}_{1-\alpha,1}$ and $\hat{c}_{1-\alpha,2}$ solve

$$\Pr\left(\max_{1\leq i\leq d_1} Z_i/\hat{\sigma}_i \leq \hat{c}_{1-\alpha,1},\ \max_{d_1+1\leq i\leq d_2} |Z_i/\hat{\sigma}_i| \leq \hat{c}_{1-\alpha,2}\right) = 1-\alpha\,,\quad Z \sim N(0,\hat{\Sigma})\,.$$

If $d_2 = 0$, then $\hat{c}_{1-\alpha,1}$ is the $(1-\alpha)$-quantile of $\max_{1\leq i\leq d_1} Z_i/\hat{\sigma}_i$; similarly, if $d_1 = 0$, then $\hat{c}_{2,1-\alpha}$ is the $(1-\alpha)$-quantile of $\max_{1\leq i\leq d_2} |Z_i/\hat{\sigma}_i|$.

7 Conclusion

This paper introduced a framework for analyzing the sensitivity of counterfactuals to parametric assumptions about the distribution of latent variables in structural models. In particular, we derived bounds on the set of counterfactuals obtained as the distribution of latent variables spans nonparametric neighborhoods of a given parametric specification while other "structural" model features are maintained. We illustrated our procedure with empirical applications to matching models and dynamic discrete choice.

References

Aguirregabiria, V. and P. Mira (2007). Sequential estimation of dynamic discrete games. *Econometrica 75*(1), 1–53.

Allen, R. and J. Rehbeck (2019). Identification with additively separable heterogeneity. *Econometrica 87*(3), 1021–1054.

Andrews, I., M. Gentzkow, and J. M. Shapiro (2017). Measuring the sensitivity of parameter estimates to estimation moments. *The Quarterly Journal of Economics 132*(4), 1553–1592.

Andrews, I., M. Gentzkow, and J. M. Shapiro (2020). On the informativeness of descriptive statistics for structural estimates. *Econometrica* (6), 2231–2258.

Armstrong, T. B. and M. Kolesár (2021). Sensitivity analysis using approximate moment condition models. *Quantitative Economics 12*(1), 77–108.

Bajari, P., C. L. Benkard, and J. Levin (2007). Estimating dynamic models of imperfect competition. *Econometrica 75*(5), 1331–1370.

Beresteanu, A., I. Molchanov, and F. Molinari (2011). Sharp identification regions in models with convex moment predictions. *Econometrica 79*(6), 1785–1821.

Berger, J. O. (1984). The robust Bayesian viewpoint (with discussion). In J. B. Kadane (Ed.), *Robustness of Bayesian analyses*. North-Holland.

Berry, S. T. (1992). Estimation of a model of entry in the airline industry. *Econometrica 60*(4), 889–917.

Berry, S. T. and P. A. Haile (2010). Nonparametric identification of multinomial choice demand models with heterogeneous consumers. Cowles Foundation Discussion Paper No. 1718.

Berry, S. T. and P. A. Haile (2014). Identification in differentiated products markets using market level data. *Econometrica 82*(5), 1749–1797.

Bonhomme, S. and M. Weidner (2021). Minimizing sensitivity to model misspecification. *Quantitative Economics*, forthcoming.

Bresnahan, T. F. and P. C. Reiss (1990). Entry in monopoly market. *The Review of Economic Studies 57*(4), 531–553.

Bresnahan, T. F. and P. C. Reiss (1991). Empirical models of discrete games. *Journal of Econometrics 48*(1), 57–81.

Chamberlain, G. and E. E. Leamer (1976). Matrix weighted averages and posterior bounds. *Journal of the Royal Statistical Society. Series B (Methodological) 38*(1), 73–84.

Chen, X., T. M. Christensen, and E. Tamer (2018). Monte carlo confidence sets for identified sets. *Econometrica 86*(6), 1965–2018.

Chen, X., E. Tamer, and A. Torgovitsky (2011). Sensitivity analysis in semiparametric likelihood models. Cowles Foundation Discussion Paper No. 1836.

Chesher, A. and A. M. Rosen (2017). Generalized instrumental variable models. *Econometrica 85*(3), 959–989.

Chiappori, P.-A., B. Salanié, and Y. Weiss (2017). Partner choice, investment in children, and the marital college premium. *American Economic Review 107*(8), 2109–2167.

Choo, E. and A. Siow (2006). Who marries whom and why. *Journal of Political Economy 114*(1), 175–201.

Christensen, T. and B. Connault (2022). Counterfactual sensitivity and robustness. arxiv preprint, arxiv:1904.00989 [econ.em].

Ciliberto, F. and E. Tamer (2009). Market structure and multiple equilibria in airline markets. *Econometrica 77*(6), 1791–1828.

Csiszár, I. (1975). *I*-divergence geometry of probability distributions and minimization problems. *The Annals of Probability 3*(1), 146–158.

Csiszár, I. and P. C. Shields (2004). Information theory and statistics: A tutorial. *Communications and Information Theory 1*(4), 417–528.

Cuturi, M. (2013). Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in Neural Information Processing Systems 26*, 2292–2300.

Dagsvik, J. K. (2000). Aggregation in matching markets. *International Economic Review 41*(1), 27–58.

Duchi, J. C. and H. Namkoong (2021). Learning models with uniform performance via distributionally robust optimization. *The Annals of Statistics 49*(3), 1378–1406.

Ekeland, I., A. Galichon, and M. Henry (2010). Optimal transportation and the falsifiability of incompletely specified economic models. *Economic Theory 42*(2), 355–374.

Fang, Z. and A. Santos (2019). Inference on directionally differentiable functions. *Review of Economic Studies 86*(1), 377–412.

Galichon, A. and M. Henry (2011). Set Identification in Models with Multiple Equilibria. *The Review of Economic Studies 78*(4), 1264–1298.

Galichon, A. and B. Salanié (2021). Cupid's invisible hand: Social surplus and identification in matching models. *The Review of Economic Studies*, forthcoming.

Giacomini, R., T. Kitagawa, and H. Uhlig (2016). Estimation under ambiguity. Technical report, UCL and Chicago.

Grieco, P. L. E. (2014). Discrete games with flexible information structures: an application to local grocery markets. *The RAND Journal of Economics 45*(2), 303–340.

Gualdani, C. and S. Sinha (2020). Partial identification in matching models for the marriage market. arxiv preprint, arxiv:1902.05610 [econ.em].

Hansen, L. P. and T. J. Sargent (2001). Robust control and model uncertainty. *The American Economic Review 91*(2), 60–66.

Hansen, L. P. and T. J. Sargent (2008). *Robustness*. Princeton.

Ho, P. (2018). Global robust bayesian analysis in large models. Technical report, Princeton.

Kalouptsidi, M., Y. Kitamura, L. Lima, and E. Souza-Rodrigues (2020). Partial identification and inference for dynamic models and counterfactuals. *NBER Working Paper No. 26761*.

Kalouptsidi, M., P. T. Scott, and E. Souza-Rodrigues (2021). Identification of counterfactuals in dynamic discrete choice models. *Quantitative Economics 12*(2), 351–403.

Kitamura, Y., T. Otsu, and K. Evdokimov (2013). Robustness, infinitesimal neighborhoods, and moment restrictions. *Econometrica 81*(3), 1185–1201.

Kline, B. and E. Tamer (2016). Bayesian inference in a class of partially identified models. *Quantitative Economics 7*(2), 329–366.

Komunjer, I. and G. Ragusa (2016). Existence and characterization of conditional density projections. *Econometric Theory 32*(4), 947–987.

Lafférs, L. (2019). Bounding average treatment effects using linear programming. *Empirical Economics 57*(3), 727–767.

Leamer, E. E. (1982). Sets of posterior means with bounded variance priors. *Econometrica 50*(3), 725–736.

Leamer, E. E. (1985). Sensitivity analyses would help. *The American Economic Review 75*(3), 308–313.

Li, L. (2018). Identification of structural and counterfactual parameters in a large class

of structural econometric models. Technical report, Pennsylvania State University.

Liese, F. and I. Vajda (1987). *Convex statistical distances*. Teubner, Leipzig.

Maccheroni, F., M. Marinacci, and A. Rustichini (2006). Ambiguity aversion, robustness, and the variational representation of preferences. *Econometrica 74*(6), 1447–1498.

Manski, C. F. (2007). Partial identification of counterfactual choice probabilities. *International Economic Review 48*(4), 1393–1410.

Manski, C. F. (2014). Identification of income–leisure preferences and evaluation of income tax policy. *Quantitative Economics 5*(1), 145–174.

McFadden, D. L. (1974). *Conditional Logit Analysis of Qualitative Choice Behavior*, pp. 105–142. Frontiers in Econometrics. Academic Press.

McFadden, D. L. (1978). *Modelling the Choice of Residential Location*, pp. 75–96. Spatial Interaction Theory and Planning Models. North Holland: Amsterdam.

Milgrom, P. and I. Segal (2002). Envelope theorems for arbitrary choice sets. *Econometrica 70*(2), 583–601.

Mukhin, Y. (2018). Sensitivity of regular estimators. arxiv preprint, arxiv:1805.08883 [econ.em].

Norets, A. and X. Tang (2014). Semiparametric inference in dynamic binary choice models. *The Review of Economic Studies 81*(3), 1229–1262.

Owen, A. B. (2017). A randomized Halton algorithm in R. arxiv:1706.02808 [stat.co].

Qiao, Y. and N. Minematsu (2010). A study on invariance of f-divergence and its application to speech recognition. *IEEE Transactions on Signal Processing 58*(7), 3884–3890.

Rockafellar, R. T. (1970). *Convex Analysis*. Princeton University Press.

Rust, J. (1987). Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica 55*(5), 999–1033.

Schennach, S. M. (2014). Entropic latent variable integration via simulation. *Econometrica 82*(1), 345–385.

Shapiro, A. (1990). On concepts of directional differentiability. *Journal of Optimization Theory and Applications 66*(3), 477–487.

Shapiro, A. (2017). Distributionally robust stochastic programming. *SIAM Journal on Optimization 27*(4), 2258–2275.

Su, C.-L. and K. L. Judd (2012). Constrained optimization approaches to estimation of structural models. *Econometrica 80*(5), 2213–2230.

Tamer, E. (2003). Incomplete simultaneous discrete response model with multiple equilibria. *The Review of Economic Studies 70*(1), 147–165.

Tamer, E. (2015). Sensitivity analysis in some econometric models. Cowles Lecture, Econometric Society World Congress, August 17-21, 2015, Montréal.

Tebaldi, P., A. Torgovitsky, and H. Yang (2019). Nonparametric estimates of demand in the california health insurance exchange. NBER working paper.

Torgovitsky, A. (2019a). Nonparametric inference on state dependence in unemployment. *Econometrica 87*(5), 1475–1505.

Torgovitsky, A. (2019b). Partial identification by extending subdistributions. *Quantitative Economics 10*(1), 105–144.

Online Appendix to "Counterfactual Sensitivity and Robustness"

Timothy Christensen   Benjamin Connault

This supplement presents extensions of our methodology in Appendix A, additional results on nonparametric bounds on counterfactuals in Appendix B, connections with local approaches to sensitivity analysis in Appendix C, additional details on the empirical applications in Appendix D, and proofs of results from the main text in Appendix E.

A   Extensions

This appendix presents three extensions of our methodology. Proofs of all results in this appendix are presented in Appendix G.7 of our working paper version Christensen and Connault (2022).

A.1   Group Invariance

In certain settings it can be attractive to impose shape restrictions on F such as symmetry, exchangeability, or, more generally, invariance to a finite group of transforms. For instance, imposing exchangeability of F in discrete choice modeling ensures that alternatives' choice probabilities depend on their deterministic components of utility but not their labeling. These shape restrictions can be easily imposed whenever F_* is invariant.

Formally, let J denote the number of elements of U and let Π be a finite commutative group of transforms on $\mathbb{R}^J$—see, e.g., Section 1.4 of Lehmann and Casella (1998). We say that a distribution F of U is Π-*invariant* if $\varpi U \sim F$ for all $\varpi \in \Pi$.

Example A.1 (Symmetry)   Central symmetry corresponds to $\Pi = \{I, -I\}$ for I the identity matrix. Sign symmetry corresponds to taking Π to be the collection of all 2^J diagonal matrices with ± 1 in each diagonal entry.   □

Example A.2 (Exchangeability)   Let Π_J denote the group of all $J!$ permutation matrices of dimension J. Full exchangeability (permutation invariance) corresponds to $\Pi = \Pi_J$. Cyclic exchangeability (rotation invariance) corresponds to $\Pi = \Pi_J^c$ where Π_J^c is the collection of all J cyclic permutation matrices of dimension J ($\Pi_J^c = \Pi_J$ when $J = 2$ and is a strict subset otherwise). When $J \geq 3$, dihedral exchangeability

(rotation and reflection invariance) corresponds to taking Π to be the set of all $2J$ permutation matrices representing rotations and reflections of $\{1, \ldots, J\}$. These types of exchangeability ensure the elements of U are identically distributed, but they have different implications for the joint distribution of the elements of U. For instance, the distribution of (U_i, U_j) for $i \neq j$ depends on $i - j$ and $|i - j|$ under cyclic and dihedral exchangeability, but is independent of (i, j) under full exchangeability. $\qquad\square$

Let $\mathcal{N}_\delta^\Pi = \{F \in \mathcal{N}_\delta : F \text{ is } \Pi\text{-invariant}\}$. We are interested in

$$\underline{\kappa}_\delta^\Pi := \inf_{\theta \in \Theta, F \in \mathcal{N}_\delta^\Pi} \mathbb{E}^F[k(U, \theta, \gamma_0)] \quad \text{subject to (1)}, \tag{32}$$

and $\overline{\kappa}_\delta^\Pi$ defined as the analogous supremum. One may write $\underline{\kappa}_\delta^\Pi$ and $\overline{\kappa}_\delta^\Pi$ as the value of two optimization problems in which criterion functions $\underline{K}_\delta^\Pi(\theta; \gamma_0, P_0)$ and $\overline{K}_\delta^\Pi(\theta; \gamma_0, P_0)$ are optimized with respect to θ. For a generic (θ, γ, P), define

$$\underline{K}_\delta^\Pi(\theta; \gamma, P) = \inf_{F \in \mathcal{N}_\delta^\Pi} \mathbb{E}^F[k(U, \theta, \gamma)] \quad \text{subject to (1) holding at } (\theta, \gamma, P), \tag{33}$$

and define $\overline{K}_\delta^\Pi(\theta; \gamma, P)$ as the analogous supremum. These criterions have dual representations as finite-dimensional convex programs when F_* is Π-invariant. Define

$$k^\Pi(U, \theta, \gamma) = \frac{1}{|\Pi|} \sum_{\varpi \in \Pi} k(\varpi U, \theta, \gamma), \quad g_j^\Pi(U, \theta, \gamma) = \frac{1}{|\Pi|} \sum_{\varpi \in \Pi} g_j(\varpi U, \theta, \gamma), \quad j = 1, 2, 3, 4,$$

where $|\Pi|$ denotes the cardinality of Π, and let $g^\Pi = (g_1^\Pi, g_2^\Pi, g_3^\Pi, g_4^\Pi)$.

Proposition A.1 *Suppose that Assumption Φ holds and F_* is Π-invariant. Then*

$$\underline{K}_\delta^\Pi(\theta; \gamma, P) = \sup_{\eta > 0, \zeta \in \mathbb{R}, \lambda \in \Lambda} -\eta \mathbb{E}^{F_*}\left[\phi^\star\left(\frac{k^\Pi(U, \theta, \gamma) + \zeta + \lambda' g^\Pi(U, \theta, \gamma)}{-\eta}\right)\right] - \eta\delta - \zeta - \lambda_{12}' P, \tag{34}$$

$$\overline{K}_\delta^\Pi(\theta; \gamma, P) = \inf_{\eta > 0, \zeta \in \mathbb{R}, \lambda \in \Lambda} \eta \mathbb{E}^{F_*}\left[\phi^\star\left(\frac{k^\Pi(U, \theta, \gamma) - \zeta - \lambda' g^\Pi(U, \theta, \gamma)}{\eta}\right)\right] + \eta\delta + \zeta + \lambda_{12}' P. \tag{35}$$

Moreover, the value of problem (34) is $+\infty$ (equivalently, the value of problem (35) is $-\infty$) if and only if there is no distribution in $\mathcal{N}_\delta^\Pi$ under which (1) holds at (θ, γ, P).

Remark A.1 *If F is Π-invariant and satisfies (1), then it must also satisfy (1) under all $|\Pi|$ transformations of the elements of U. Therefore, in effect there are a total of*

$|\Pi| \times d$ *moment conditions imposed in the inner optimization, namely*

$$\mathbb{E}^F[g_1(\varpi U, \theta, \gamma_0)] \leq P_{10}, \quad \mathbb{E}^F[g_2(\varpi U, \theta, \gamma_0)] = P_{20},$$
$$\mathbb{E}^F[g_3(\varpi U, \theta, \gamma_0)] \leq 0, \qquad \mathbb{E}^F[g_4(\varpi U, \theta, \gamma_0)] = 0, \qquad \textit{for all } \varpi \in \Pi. \qquad (36)$$

In principle one could form a criterion by including all $|\Pi| \times d$ moments. By Π-invariance of F_* and convexity of the objective, the multipliers on the moments $g(\varpi U, \theta, \gamma)$ will be identical across all $\varpi \in \Pi$. It therefore suffices to form the criterion using only the d averaged moments g^Π rather than the full set of $|\Pi| \times d$ moments, thereby reducing the dimension of the inner optimization by a factor of $|\Pi|$.

Remark A.2 *When Monte Carlo integration is used to compute expectations, taking a sample from F_* and then concatenating the sample across each of its $|\Pi|$ transformations ensures the empirical distribution of the random draws is Π-invariant.*

A.2   Conditional Moment Models

Consider the conditional moment model

$$\mathbb{E}^F[g_1(U, X, \theta, \gamma_0)|X = x] \leq P_{10,x}, \quad \mathbb{E}^F[g_2(U, X, \theta, \gamma_0)|X = x] = P_{20,x},$$
$$\mathbb{E}^F[g_3(U, X, \theta, \gamma_0)|X = x] \leq 0, \qquad \mathbb{E}^F[g_4(U, X, \theta, \gamma_0)|X = x] = 0, \qquad \text{for all } x \in \mathcal{X}$$
$$(37)$$

where $\mathcal{X}$ is a finite set, and a counterfactual[24]

$$\kappa = \sum_{x \in \mathcal{X}} \mathbb{E}^F[k(U, X, \theta, \gamma_0)|X = x]. \qquad (38)$$

Suppose the researcher assumes $U|X = x \sim F_*$ for each x. We wish to relax this assumption and allow each conditional distribution of U given $X = x$, say F_x, to vary in a neighborhood $\mathcal{N}_{\delta_x}$ of F_*. In doing so, we are allowing the conditional distributions F_x to vary with x, and therefore relaxing independence of U and X.[25]

We assume each $\mathcal{N}_\delta$ is defined by the same ϕ to simplify the exposition, but we allow

---

[24]Note κ can be the expected value at a particular x_0 if $k(U, x, \theta, \gamma_0) = 0$ for $x \neq x_0$. More generally, κ can be a weighted average by incorporating the weighting into the definition of $k(u, x, \theta, \gamma_0)$.

[25]The case with U independent of X is subsumed in (1) by stacking the moment functions and reduced-form parameters by values of the conditioning variable, as in Examples 2.1–2.3.

the neighborhood size to vary with x. Let $\boldsymbol{\delta} = (\delta_x)_{x \in \mathcal{X}}$. We are interested in

$$\underline{\kappa}_{\boldsymbol{\delta}} := \inf_{\theta \in \Theta, (F_x \in \mathcal{N}_{\delta_x})_{x \in \mathcal{X}}} \sum_{x \in \mathcal{X}} \mathbb{E}^{F_x}[k(U, x, \theta, \gamma_0)] \quad \text{subject to (37),} \tag{39}$$

and $\overline{\kappa}_{\boldsymbol{\delta}}$ defined as the analogous supremum. One may write $\underline{\kappa}_{\boldsymbol{\delta}}$ and $\overline{\kappa}_{\boldsymbol{\delta}}$ as the value of two optimization problems where $\underline{K}_{\boldsymbol{\delta}}(\theta; \gamma_0, P_0)$ and $\overline{K}_{\boldsymbol{\delta}}(\theta; \gamma_0, P_0)$ are optimized with respect to θ. Let $P = (P_x)_{x \in \mathcal{X}}$ where $P_x = (P_{1,x}, P_{2,x})$ is partitioned conformably with g_1 and g_2. For a generic (θ, γ, P), define

$$\underline{K}_{\boldsymbol{\delta}}(\theta; \gamma, P) = \inf_{(F_x \in \mathcal{N}_{\delta_x})_{x \in \mathcal{X}}} \sum_{x \in \mathcal{X}} \mathbb{E}^{F_x}[k(U, x, \theta, \gamma_0)] \quad \text{subject to (37) holding at } (\theta, \gamma, P),$$

and define $\overline{K}_{\boldsymbol{\delta}}(\theta; \gamma, P)$ as the analogous supremum. These criterion functions have dual forms analogous to Proposition 2.1. Let $g(\cdot, x, \theta, \gamma) = (g_1(\cdot, x, \theta, \gamma), \ldots, g_4(\cdot, x, \theta, \gamma))$. Recall $d = \sum_{i=1}^{4} d_i$ where d_i is the dimension of g_i, and $\Lambda = \mathbb{R}_+^{d_1} \times \mathbb{R}^{d_2} \times \mathbb{R}_+^{d_3} \times \mathbb{R}^{d_4}$. Let $\lambda_{12,x}$ denote the first $d_1 + d_2$ elements of $\lambda_x \in \Lambda$.

Assumption Φ-conditional *(i) $\phi \in \Phi_0$.*

(ii) $k(\,\cdot\,, x, \theta, \gamma)$ and each entry of $g(\,\cdot\,, x, \theta, \gamma)$ belong to $\mathcal{E}$ for each $(\theta, \gamma, x) \in \Theta \times \Gamma \times \mathcal{X}$.

Proposition A.2 *Suppose that Assumption Φ-conditional holds. Then*

$$\underline{K}_{\boldsymbol{\delta}}(\theta; \gamma, P) \tag{40}$$
$$= \sup_{(\eta_x > 0, \zeta_x \in \mathbb{R}, \lambda_x \in \Lambda)_{x \in \mathcal{X}}} \sum_{x \in \mathcal{X}} \left(-\eta_x \mathbb{E}^{F_*} \left[\phi^{\star} \left(\frac{k(U, x, \theta, \gamma) + \zeta_x + \lambda'_x g(U, x, \theta, \gamma)}{-\eta_x} \right) \right] - \eta_x \delta_x - \zeta_x - \lambda'_{12,x} P_x \right),$$

$$\overline{K}_{\boldsymbol{\delta}}(\theta; \gamma, P) \tag{41}$$
$$= \inf_{(\eta_x > 0, \zeta_x \in \mathbb{R}, \lambda_x \in \Lambda)_{x \in \mathcal{X}}} \sum_{x \in \mathcal{X}} \left(\eta_x \mathbb{E}^{F_*} \left[\phi^{\star} \left(\frac{k(U, x, \theta, \gamma) - \zeta_x - \lambda'_x g(U, x, \theta, \gamma)}{\eta_x} \right) \right] + \eta_x \delta_x + \zeta_x + \lambda'_{12,x} P_x \right).$$

Moreover, the value of (40) is $+\infty$ (equivalently, the value of (41) is $-\infty$) if and only if for some $x \in \mathcal{X}$ there is no distribution in $\mathcal{N}_{\delta_x}$ under which (37) holds at (θ, γ, P).

As before, estimators $\hat{\underline{\kappa}}_{\boldsymbol{\delta}}$ and $\hat{\overline{\kappa}}_{\boldsymbol{\delta}}$ of $\underline{\kappa}_{\boldsymbol{\delta}}$ and $\overline{\kappa}_{\boldsymbol{\delta}}$ are computed by optimizing sample criterions with respect to θ. Let $\hat{P} = (\hat{P}_x)_{x \in \mathcal{X}}$. The sample criterions are

$$\hat{\underline{K}}_{\boldsymbol{\delta}}(\theta) = \left[\begin{array}{l} \underline{K}_{\boldsymbol{\delta}}(\theta; \hat{\gamma}, \hat{P}) \\ +\infty \end{array} \right., \quad \hat{\overline{K}}_{\boldsymbol{\delta}}(\theta) = \left[\begin{array}{ll} \overline{K}_{\boldsymbol{\delta}}(\theta; \hat{\gamma}, \hat{P}) & \text{if } \Delta_x(\theta; \hat{\gamma}, \hat{P}_x) < \delta_x \text{ for each } x \in \mathcal{X}, \\ -\infty & \text{if } \Delta_x(\theta; \hat{\gamma}, \hat{P}_x) \geq \delta_x \text{ for some } x \in \mathcal{X}, \end{array} \right.$$

4

where $\overline{K}_{\boldsymbol{\delta}}(\theta; \hat{\gamma}, \hat{P})$ and $\overline{K}_{\boldsymbol{\delta}}(\theta; \hat{\gamma}, \hat{P})$ denote the programs in Proposition A.2 evaluated at $(\hat{\gamma}, \hat{P})$, and

$$\Delta_x(\theta; \hat{\gamma}, \hat{P}_x) = \sup_{\zeta_x \in \mathbb{R}, \lambda_x \in \Lambda} -\mathbb{E}^{F*}\Big[\phi^\star(-\zeta_x - \lambda_x' g(U, x, \theta, \hat{\gamma}))\Big] - \zeta_x - \lambda_{12,x}' \hat{P}_x.$$

A.3 Non-separable Models

Consider the model

$$\begin{aligned} \mathbb{E}^H[\tilde{g}_1(U, X, \theta, \tilde{\gamma}_0)] &\leq P_{10}, & \mathbb{E}^H[\tilde{g}_2(U, X, \theta, \tilde{\gamma}_0)] &= P_{20}, \\ \mathbb{E}^H[\tilde{g}_3(U, X, \theta, \tilde{\gamma}_0)] &\leq 0, & \mathbb{E}^H[\tilde{g}_4(U, X, \theta, \tilde{\gamma}_0)] &= 0, \end{aligned} \tag{42}$$

and counterfactual

$$\kappa = \mathbb{E}^H[\tilde{k}(U, X, \theta, \tilde{\gamma}_0)], \tag{43}$$

where the expectation is with respect to the distribution H of (U, X) and X takes values in a finite set $\mathcal{X}$. Suppose the researcher assumes $U|X = x \sim F_*$ for each x. We wish to relax this assumption and allow the conditional distribution of U given $X = x$, say F_x, to vary in a neighborhood $\mathcal{N}_{\delta_x}$ of F_*.

Write $H(u, x) = q_{0,x} \cdot F_x(u)$ where $q_{0,x} = \Pr(X = x)$. The vector $q_0 = (q_{0,x})_{x \in \mathcal{X}}$ can be consistently estimated from data on X. Let $\gamma_0 = (\tilde{\gamma}_0, q_0)$. Define $g_1(U, x, \theta, \gamma_0) = q_{0,x} \cdot \tilde{g}_1(U, x, \theta, \tilde{\gamma}_0)$ and similarly for g_2, g_3, g_4, and k. The model (42) and counterfactual (43) can then be written as

$$\begin{aligned} \sum_x \mathbb{E}^{F_x}[g_1(U, x, \theta, \gamma_0)] &\leq P_{10}, & \sum_x \mathbb{E}^{F_x}[g_2(U, x, \theta, \gamma_0)] &= P_{20}, \\ \sum_x \mathbb{E}^{F_x}[g_3(U, x, \theta, \gamma_0)] &\leq 0, & \sum_x \mathbb{E}^{F_x}[g_4(U, x, \theta, \gamma_0)] &= 0, \end{aligned} \tag{44}$$

and $\kappa = \sum_x \mathbb{E}^{F_x}[k(U, x, \theta, \gamma_0)]$. We again assume each $\mathcal{N}_\delta$ is defined by the same ϕ function, but allow the neighborhood size to vary with x. Let $\boldsymbol{\delta} = (\delta_x)_{x \in \mathcal{X}}$. We are interested in

$$\underline{\kappa}_{\boldsymbol{\delta}} := \inf_{\theta \in \Theta, (F_x \in \mathcal{N}_{\delta_x})_{x \in \mathcal{X}}} \sum_x \mathbb{E}^{F_x}[k(U, x, \theta, \gamma_0)] \quad \text{subject to (44)},$$

and $\overline{\kappa}_{\boldsymbol{\delta}}$ defined as the analogous supremum. One may write $\underline{\kappa}_{\boldsymbol{\delta}}$ and $\overline{\kappa}_{\boldsymbol{\delta}}$ as the value of two optimization problems where criterion functions $\underline{K}_{\boldsymbol{\delta}}(\theta; \gamma_0, P_0)$ and $\overline{K}_{\boldsymbol{\delta}}(\theta; \gamma_0, P_0)$ are

optimized with respect to θ. For a generic (θ, γ, P), define

$$\underline{K}_{\boldsymbol{\delta}}(\theta; \gamma, P) = \inf_{(F_x \in \mathcal{N}_{\delta_x})_{x \in \mathcal{X}}} \sum_x \mathbb{E}^{F_x}[k(U, x, \theta, \gamma_0)] \quad \text{s.t. (44) holding at } (\theta, \gamma, P),$$

and define $\overline{K}_{\boldsymbol{\delta}}(\theta; \gamma, P)$ as the analogous supremum. These criterion functions have dual forms analogous to Proposition 2.1. Let $g(\cdot, x, \theta, \gamma) = (g_1(\cdot, x, \theta, \gamma), \ldots, g_4(\cdot, x, \theta, \gamma))$. The remaining notation the same as Proposition 2.1.

Proposition A.3 *Suppose that Assumption Φ-conditional holds. Then*

$$\underline{K}_{\boldsymbol{\delta}}(\theta; \gamma, P) \tag{45}$$
$$= \sup_{(\eta_x > 0, \zeta_x \in \mathbb{R})_{x \in \mathcal{X}}, \lambda \in \Lambda} \sum_x \left(-\eta_x \mathbb{E}^{F_*} \left[\phi^\star \left(\tfrac{k(U, x, \theta, \gamma) + \zeta_x + \lambda' g(U, x, \theta, \gamma)}{-\eta_x} \right) \right] - \eta_x \delta_x - \zeta_x - \lambda_{12}' P \right),$$

$$\overline{K}_{\boldsymbol{\delta}}(\theta; \gamma, P) \tag{46}$$
$$= \inf_{(\eta_x > 0, \zeta_x \in \mathbb{R})_{x \in \mathcal{X}}, \lambda \in \Lambda} \sum_x \left(\eta_x \mathbb{E}^{F_*} \left[\phi^\star \left(\tfrac{k(U, x, \theta, \gamma) - \zeta_x - \lambda' g(U, x, \theta, \gamma)}{\eta_x} \right) \right] + \eta_x \delta_x + \zeta_x + \lambda_{12}' P \right).$$

Moreover, the value of (45) is $+\infty$ (equivalently, the value of (46) is $-\infty$) if and only if there is no $H(u, x) = q_{0, x} \cdot F_x(u)$ with $F_x \in \mathcal{N}_{\delta_x}$ under which (42) holds at (θ, γ, P).

As before, estimators $\hat{\underline{\kappa}}_{\boldsymbol{\delta}}$ and $\hat{\overline{\kappa}}_{\boldsymbol{\delta}}$ of $\underline{\kappa}_{\boldsymbol{\delta}}$ and $\overline{\kappa}_{\boldsymbol{\delta}}$ are computed by optimizing sample criterion functions with respect to θ. The sample criterion functions are

$$\hat{\underline{K}}_{\boldsymbol{\delta}}(\theta) = \begin{bmatrix} \underline{K}_{\boldsymbol{\delta}}(\theta; \hat{\gamma}, \hat{P}) \\ +\infty \end{bmatrix}, \quad \hat{\overline{K}}_{\boldsymbol{\delta}}(\theta) = \begin{bmatrix} \overline{K}_{\boldsymbol{\delta}}(\theta; \hat{\gamma}, \hat{P}) & \text{if } \Delta_{nonsep}(\theta; \hat{\gamma}, \hat{P}) < 0 \\ -\infty & \text{if } \Delta_{nonsep}(\theta; \hat{\gamma}, \hat{P}) \geq 0, \end{bmatrix}$$

where $\overline{K}_{\boldsymbol{\delta}}(\theta; \hat{\gamma}, \hat{P})$ and $\overline{K}_{\boldsymbol{\delta}}(\theta; \hat{\gamma}, \hat{P})$ denote the programs in Proposition A.3 evaluated at $(\hat{\gamma}, \hat{P})$ with $\hat{\gamma} = (\hat{\tilde{\gamma}}, \hat{q})$ for estimators $\hat{\tilde{\gamma}}$ of $\tilde{\gamma}$ and $\hat{q}$ of q_0, and

$$\Delta_{nonsep}(\theta; \gamma, P)$$
$$= \sup_{\substack{(\eta_x \geq 0, \zeta_x \in \mathbb{R})_{x \in \mathcal{X}}, \lambda \in \Lambda \\ \sum_{x \in \mathcal{X}} \eta_x \leq 1}} \left(-\sum_{x \in \mathcal{X}} \mathbb{E}^{F_*} \left[(\eta_x \phi)^\star (-\zeta_x - \lambda_x' g(U, x, \theta, \gamma)) \right] - \eta_x \delta_x - \zeta_x \right) - \lambda_{12}' P.$$

By similar arguments to Appendix G.3 of Christensen and Connault (2022), $\Delta_{nonsep}(\theta; \gamma, P)$ may be shown to be the dual of

$$\inf_{t \in \mathbb{R}, (F_x)_{x \in \mathcal{X}}} t \quad \text{s.t.} \quad D_\phi(F_x \| F_*) \leq \delta_x + t \text{ for each } x \in \mathcal{X} \text{ and (44) holding at } (\theta, \gamma, P).$$

6

Therefore, if there exists F_x with $D_\phi(F_x\|F_*) < \delta_x$ for each x such that (44) holds at (θ, γ, P), then $\Delta_{nonsep}(\theta; \gamma, P) < 0$.

B   Additional Results on Nonparametric Bounds

This appendix presents further details to supplement Section 2.5. Proofs of all results in this appendix are presented in Appendix G.8 of Christensen and Connault (2022). Our first result concerns the behavior of $\underline{\kappa}_\delta$ and $\overline{\kappa}_\delta$ as the neighborhood size δ becomes large. Recall $\mathcal{N}_\infty = \{F : D_\phi(F\|F_*) < \infty\}$. Let

$$\mathcal{K}_\infty = \{\mathbb{E}^F[k(U, \theta, \gamma_0)] : (1) \text{ holds at } (\theta, \gamma_0, P_0) \text{ for some } \theta \in \Theta,\, F \in \mathcal{N}_\infty\}.$$

Lemma B.1 *Suppose that Assumption Φ holds. Then*

$$\lim_{\delta \to \infty} \underline{\kappa}_\delta = \inf \mathcal{K}_\infty\,, \qquad\qquad \lim_{\delta \to \infty} \overline{\kappa}_\delta = \sup \mathcal{K}_\infty\,.$$

Next, we characterize bounds on $\mathcal{K}_\infty$ using profiled optimization problems and derive their dual forms. Define

$$\underline{K}_\infty(\theta; \gamma_0, P_0) = \inf_{F \in \mathcal{N}_\infty} \mathbb{E}^F[k(U, \theta, \gamma_0)] \quad \text{subject to (1) holding at } (\theta, F)\,, \qquad (47)$$

and let $\overline{K}_\infty(\theta; \gamma_0, P_0)$ denote the analogous supremum. By definition, we have

$$\inf \mathcal{K}_\infty = \inf_{\theta \in \Theta} \underline{K}_\infty(\theta; \gamma_0, P_0)\,, \qquad \sup \mathcal{K}_\infty = \sup_{\theta \in \Theta} \overline{K}_\infty(\theta; \gamma_0, P_0)\,.$$

Let F_*-ess inf and F_*-ess sup denote the F_*-essential infimum and supremum, respectively.

Lemma B.2 *Suppose that Assumption Φ holds and Condition S holds at (θ, γ, P). Then*

$$\underline{K}_\infty(\theta; \gamma, P) = \sup_{\lambda \in \Lambda: F_*\text{-ess inf}(k(\cdot, \theta, \gamma) + \lambda' g(\cdot, \theta, \gamma)) > -\infty} (F_*\text{-ess inf}(k(\cdot, \theta, \gamma) + \lambda' g(\cdot, \theta, \gamma)) - \lambda'_{12}P)\,,$$

$$\overline{K}_\infty(\theta; \gamma, P) = \inf_{\lambda \in \Lambda: F_*\text{-ess sup}(k(\cdot, \theta, \gamma) - \lambda' g(\cdot, \theta, \gamma)) < +\infty} (F_*\text{-ess sup}(k(\cdot, \theta, \gamma) - \lambda' g(\cdot, \theta, \gamma)) + \lambda'_{12}P)\,.$$

We now derive analogous dual representations for the criterion functions $\underline{K}_{np}$ and $\overline{K}_{np}$ from Section 2.5 (see display (18)). We require a slightly different constraint qualification:

Definition B.1 Condition S_{np} *holds at* (θ, γ, P) *if* $\vec{P} \in \mathrm{ri}(\{\mathbb{E}^F[g(U, \theta, \gamma)] : F \in \mathcal{F}_\theta\} + \mathcal{C})$.

If F_* and μ are mutually absolutely continuous, then Condition S_{np} is equivalent to Condition S from Section 2.5 (see Lemma E.1).

Lemma B.3 *Suppose that Condition S_{np} holds at (θ, γ, P) and k is μ-essentially bounded. Then*

$$\underline{K}_{np}(\theta; \gamma, P) = \sup_{\lambda \in \Lambda : \mu\text{-ess inf}(k(\cdot, \theta, \gamma) + \lambda' g(\cdot, \theta, \gamma)) > -\infty} (\mu\text{-ess inf}(k(\cdot, \theta, \gamma) + \lambda' g(\cdot, \theta, \gamma)) - \lambda_{12}' P),$$

$$\overline{K}_{np}(\theta; \gamma, P) = \inf_{\lambda \in \Lambda : \mu\text{-ess sup}(k(\cdot, \theta, \gamma) - \lambda' g(\cdot, \theta, \gamma)) < +\infty} (\mu\text{-ess sup}(k(\cdot, \theta, \gamma) - \lambda' g(\cdot, \theta, \gamma)) + \lambda_{12}' P).$$

C Local Sensitivity

In this appendix, we first introduce a measure of local sensitivity of the counterfactual with respect to F. We then contrast our approach with recent work on local sensitivity.

C.1 Measure of Local Sensitivity

Our *measure of local sensitivity* of the counterfactual κ with respect to F at F_* is

$$s = \lim_{\delta \downarrow 0} \frac{(\overline{\kappa}_\delta - \underline{\kappa}_\delta)^2}{4\delta}.$$

If s is finite, then under the regularity conditions below

$$\underline{\kappa}_\delta = \kappa_* - \sqrt{\delta s} + o(\sqrt{\delta}), \quad \overline{\kappa}_\delta = \kappa_* + \sqrt{\delta s} + o(\sqrt{\delta}), \quad \text{as } \delta \downarrow 0,$$

where $\kappa_* = \mathbb{E}^{F_*}[k(U, \theta_*, \gamma_0)]$ and θ_* solves (1) under F_*.

To draw connections with the local sensitivity literature, we restrict attention to moment equality models and impose (standard) regularity conditions. These conditions allow us to characterize s very tractably via an influence function representation, which leads to a simple estimator $\hat{s}$ of s. Assume that under F_* the moment conditions (1b) and (1d) point identify a structural parameter $\theta_* \in \mathrm{int}(\Theta)$, where we again assume Θ is compact. With some abuse of notation, let

$$g(u, \theta, \gamma, P_2) = \begin{bmatrix} g_2(u, \theta, \gamma) - P_2 \\ g_4(u, \theta, \gamma) \end{bmatrix},$$

8

$g_*(u) = g(u, \theta_*, \gamma_0, P_{20})$, and $k_*(u) = k(u, \theta_*, \gamma_0)$. Let $\mathbb{E}^{F_*}[g(U, \theta, \gamma_0, P_{20})]$ and $\mathbb{E}^{F_*}[k(U, \theta, \gamma_0)]$ be continuously differentiable with respect to θ at θ_*, $G = \frac{\partial}{\partial \theta'} \mathbb{E}^{F_*}[g(U, \theta, \gamma_0, P_{20})]\big|_{\theta=\theta_*}$ have full rank, $V = \mathbb{E}^{F_*}[g_*(U)g_*(U)']$ be finite and positive definite, $\mathbb{E}^{F_*}[k(U, \theta_*, \gamma_0)^2]$ be finite, and $k(\cdot, \theta, \gamma_0)$ and $g(\cdot, \theta, \gamma_0, P_{20})$ be $L^2(F_*)$-continuous in θ at θ_*.

Define the *influence function* of the counterfactual κ with respect to F at F_* as

$$\iota(u) = \mathbb{M}k_*(u) - J'(G'V^{-1}G)^{-1}G'V^{-1}g_*(u),$$

where $\mathbb{M}k_*(u) = k_*(u) - \kappa_* - \mathbb{E}^{F_*}[k_*(U)g_*(U)'](V^{-1} - V^{-1}G(G'V^{-1}G)^{-1}G'V^{-1})g_*(u)$ and $J = \frac{\partial}{\partial \theta} \mathbb{E}^{F_*}[k(U, \theta, \gamma_0)]\big|_{\theta=\theta_*}$. The following theorem relates s and ι. We restrict attention to neighborhoods characterized by χ^2 divergence. Other ϕ-divergences are locally equivalent to χ^2 divergence, so this restriction entails no great loss of generality.[26]

Theorem C.1 *Suppose that the above GMM-type regularity conditions hold and neighborhoods are defined using χ^2 divergence. Then $s = 2\mathbb{E}^{F_*}[\iota(U)^2]$.*

The proof of Theorem C.1 is presented in Appendix G.9 of our working paper version Christensen and Connault (2022). In addition to reporting an estimated counterfactual $\hat{\kappa} = \mathbb{E}^{F_*}[k(U, \hat{\theta}, \hat{\gamma})]$, researchers could also report an estimate of its local sensitivity to F:

$$\hat{s} = 2\mathbb{E}^{F_*}[(\hat{k}(U) - \hat{\kappa})^2] + 2\hat{Q}'\hat{V}\hat{Q} - 4\mathbb{E}^{F_*}[\hat{g}(U)(\hat{k}(U) - \hat{\kappa})]'\hat{Q},$$

where $\hat{k}(u) = k(u, \hat{\theta}, \hat{\gamma})$, $\hat{g}(u) = g(u, \hat{\theta}, \hat{\gamma}, \hat{P}_2)$, $\hat{V} = \mathbb{E}^{F_*}[\hat{g}(U)\hat{g}(U)']$, and

$$\hat{Q}' = \mathbb{E}^{F_*}[\hat{k}(U)\hat{g}(U)'](\hat{V}^{-1} - \hat{V}^{-1}\hat{G}(\hat{G}'\hat{V}^{-1}\hat{G})^{-1}\hat{G}'\hat{V}^{-1}) + \hat{J}'(\hat{G}'\hat{V}^{-1}\hat{G})^{-1}\hat{G}'\hat{V}^{-1},$$

with $\hat{G} = \frac{\partial}{\partial \theta'} \mathbb{E}^{F_*}[g(U, \theta, \hat{\gamma}, \hat{P}_2)]\big|_{\theta=\hat{\theta}}$ and $\hat{J} = \frac{\partial}{\partial \theta} \mathbb{E}^{F_*}[k(U, \theta, \hat{\gamma})]\big|_{\theta=\hat{\theta}}$. Lemma G.12 in Appendix G.9 of Christensen and Connault (2022) shows $\hat{s}$ is consistent. Bounds on counterfactuals as F varies over small neighborhoods of F_* can then be estimated using $\hat{\kappa} \pm \sqrt{\delta \hat{s}}$.

C.2 Comparison with Other Approaches

We now compare our approach with Andrews et al. (2017, 2020), henceforth AGS, and Bonhomme and Weidner (2021), henceforth BW. To simplify the comparison, we con-

---

[26]See Theorem 4.1 of Csiszár and Shields (2004). The quantity $2\mathbb{E}^{F_*}[\iota(U)^2]$ should be rescaled by a factor of $\phi''(1)$ for other ϕ divergences.

sider models characterized by moments of the form (1b) with $d_2 \geq d_\theta$ and in which there is no γ.

AGS consider a setting in which the moments (1b) are locally misspecified:

$$\mathbb{E}^{F_*}[g_2(U, \theta_*)] = P_{20} + n^{-1/2}c. \tag{48}$$

Suppose a researcher has a first-stage estimator $\hat{P}_2$, computes an estimator $\hat{\theta}$ by minimizing

$$(\mathbb{E}^{F_*}[g_2(U, \theta)] - \hat{P}_2)'\hat{W}(\mathbb{E}^{F_*}[g_2(U, \theta)] - \hat{P}_2),$$

then estimates the counterfactual using $\hat{\kappa} = \mathbb{E}^{F_*}[k(U, \hat{\theta})]$. AGS's measure of *sensitivity* of $\hat{\kappa}$ to $\hat{P}_2$ is $J'(G'WG)^{-1}G'W$, where W is the probability limit of $\hat{W}$. The first-order asymptotic bias of $\hat{\kappa}$ due to local misspecification is therefore $J'(G'WG)^{-1}G'Wc$. AGS's measure of *informativeness* of $\hat{P}_2$ for $\hat{\kappa}$ is 1, meaning that all sampling variation in $\hat{\kappa}$ is explained by sampling variation in $\hat{P}_2$. Our measure s instead characterizes "specification variation" in κ as the researcher varies F subject to the moment condition (1b).

BW consider estimation of a target parameter (κ in our context) using a reference model $\mathcal{M}_R = \{(\theta, F) \in \Theta \times \{F_*\}\}$ when the true (θ_0, F_0) possibly belongs to a larger model $\mathcal{M}_L = \{(\theta, F) \in \Theta \times \mathcal{N}_\delta\}$ with $\delta \downarrow 0$ as the sample size n increases so that $n\delta \to \tau \geq 0$. BW seek estimators of κ under $\mathcal{M}_R$ that minimize worst-case asymptotic bias or MSE over $\mathcal{M}_L$. Consider the one-step estimator

$$\hat{\kappa} = \mathbb{E}^{F_*}[k(U, \hat{\theta})] + a'(\mathbb{E}^{F_*}[g_2(U, \hat{\theta})] - \hat{P}_2),$$

where $\hat{\theta}$ is a $\sqrt{n}$-consistent estimator of θ_* and $a \in \mathbb{R}^{d_2}$ satisfies $J' = -a'G$ so that $\hat{\kappa}$ does not depend asymptotically on $\hat{\theta}$. The true counterfactual is $\kappa_0 = \mathbb{E}^{F_0}[k(U, \theta_0)]$ where $(\theta_0, F_0) \in \mathcal{M}_L$ satisfies $\mathbb{E}^{F_0}[g_2(U, \theta_0)] = P_{20}$. If $\mathcal{M}_R$ is correctly specified so that $\mathbb{E}^{F_*}[g_2(U, \theta_*)] = P_{20}$, then for any a the worst-case asymptotic bias of the one-step estimator is

$$\lim_{n\to\infty} \sup_{(\theta_0, F_0)\in\mathcal{M}_L : \mathbb{E}^{F_0}[g_2(U,\theta_0)]=P_{20}} |\sqrt{n}(\kappa_* - \kappa_0))| = \sqrt{\tau}s,$$

where s is our measure of local sensitivity.

If we allow for local misspecification of $\mathcal{M}_R$, so that $\mathbb{E}^{F_*}[g_2(U, \theta_*)] \neq P_{20}$, then the worst-case asymptotic bias of the one-step estimator is

$$\lim_{n\to\infty} \sup_{(\theta_0, F_0)\in\mathcal{M}_L : \mathbb{E}^{F_0}[g_2(U,\theta_0)]=P_{20}} \left|\sqrt{n}\left(\kappa_* - \kappa_0 + a'\left(\mathbb{E}^{F_*}[g_2(U, \theta_*)] - P_{20}\right)\right)\right| = \sqrt{\tau}s_a,$$
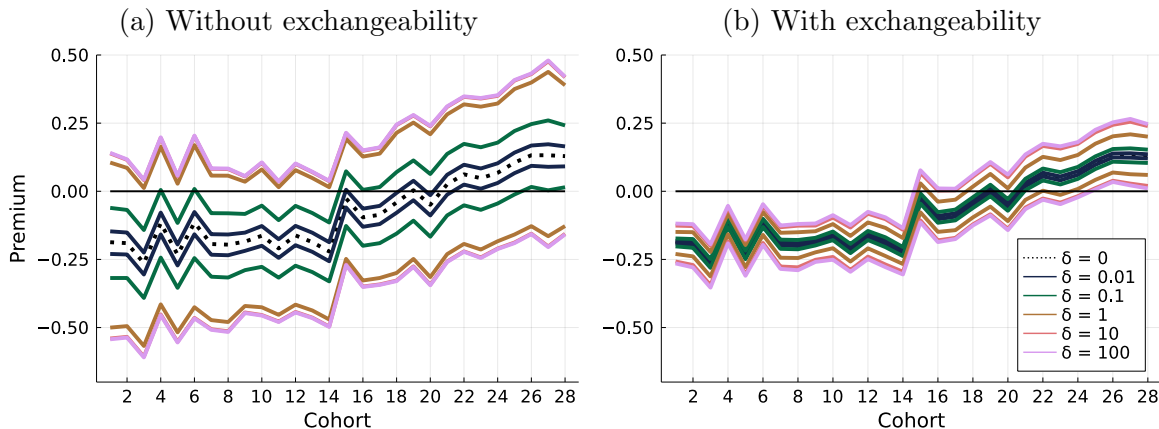
10

Figure 5: Fixed-θ bounds on the "some college" to "college graduate" premium when structural parameters are held fixed at CSW's estimates.

where s_a is our local sensitivity measure with k replaced by $k + a'g_2$.

D Additional Details for the Empirical Applications

D.1 Marital College Premium

Bootstrap Details. CSs reported Section 5.1 with $\delta > 0$ are computed using the bootstrap procedure from Section 6.2. To implement the bootstrap, we take 1,000 independent draws of $\hat{P}_2^* \sim N(\hat{P}_2, \hat{\Sigma})$ where $\hat{\Sigma}$ is CSW's estimate of the covariance matrix of $\hat{P}_2$. We compute $\hat{P}_2$ and $\hat{\Sigma}$ based on CSW's replication files.

Fixed-θ Bounds. Figure 5a plots lower and upper bounds on the "some college" to "college graduate" premium across cohorts when θ is fixed at CSW's estimates (computed under F_*) but F is allowed to vary. These bounds for large δ contain zero across each cohort, and are approximately the same width as the bounds with $\delta = 0.01$ reported in Figure 1a where both θ and F are allowed to vary. Imposing exchangeability (Figure 5b) is seen to tighten the bounds substantially, producing bounds that span negative values only for early cohorts and positive values only in the latest few cohorts.

Projection CSs. Figure 6 reports projection CSs computed using the procedure in Section 6.3. We formed 95% rectangular CSs for each cohort's P_{20} as described in Section 6.3 using CSW's estimates for $\hat{P}_2$ and their asymptotic variance estimates for $\hat{\Sigma}$. These CSs are significantly wider than the bootstrap CSs reported in Figure 1. Some conservativeness is to be expected, as these CSs project a 95% CS for P_{20} down to one
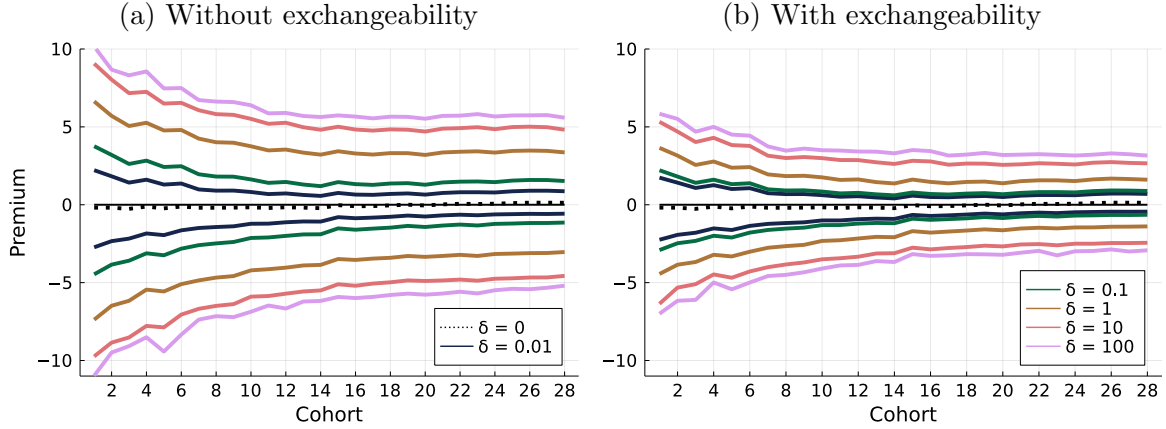
11

Figure 6: Projection 95% CSs for bounds on the "some college" to "college graduate" premium across cohorts.

Table 5: Computation times for the inner problem in the matching application

Implementation	δ				
	0.01	0.1	1	10	100
Without exchangeability	0.074	0.056	0.076	0.579	0.188
With exchangeability	0.146	0.184	0.350	0.311	0.488

Note: Times (in seconds) for solving the inner optimization problem for maximizing the premium in cohort 1 at CSW's parameter estimate θ. We use 50,000 Monte Carlo draws without exchangeability and 120,000 draws with exchangeability. All computations are performed in Julia version 1.6.4 and Knitro 12.4.0 on a 2.7GHz MacBook Pro with 16GB memory.

dimension. The relative inefficiency is especially pronounced for the earlier cohorts. Note also from Figure 6b that the projection CSs with $\delta = 0.01$ span zero across each cohort, whereas the bootstrap CSs with $\delta = 0.01$ in Figure 1b contain negative values only in some early cohorts and positive values only in later cohorts.

Computation Times. Table 5 reports times for solving the inner problem for maximizing the premium in cohort 1. This optimization problem defines the criterion function $\overline{K}_\delta(\theta; \hat{P})$. As times vary with θ, we report times at CSW's estimates. Times increase somewhat with δ, but are all under 0.6 seconds. The outer optimization times varied with cohort, δ, and implementation but were typically solved in at most a few minutes (often under 90 seconds).

Sensitivity to ϕ. Using χ^2 and L^4 divergences produced near identical bounds for $\delta = 0.01$ and 0.1. The χ^2 bounds with $\delta = 1$ and 10 were at most 10% narrower than the hybrid bounds. The L^4 bounds were 60%-70% of the width of the hybrid

bounds for $\delta = 1$, 10, and 100 across cohorts (L^4 divergence is stronger than χ^2 and hybrid divergence). The shapes of the sets were also similar to those reported for hybrid divergence. Overall, these results show that the conclusions we draw from our analysis are not sensitive to the choice of ϕ.

D.2 Welfare Analysis in a Rust Model

Bootstrap Details. Bootstrap CSs reported Section 5.2 with $\delta > 0$ are computed using the procedure from Section 6.2. We take 1,000 independent draws of $\hat{\theta}_\pi^* \sim N(\hat{\theta}_\pi, \hat{\Sigma})$ where $\hat{\theta}_\pi$ is the MLE of (RC, MC) under the i.i.d. Gumbel assumption and $\hat{\Sigma}$ is an estimate of the inverse information matrix. We then set $\hat{P}_2^*$ to be the model-implied CCPs at $\hat{\theta}_\pi^*$ under the i.i.d. Gumbel assumption.

As k depends only implicitly on u through θ, we compute $\hat{\underline{\kappa}}_\delta$ and $\hat{\overline{\kappa}}_\delta$ using the criterion functions in display (17), which is more computationally efficient than using criterions (13) and (14). The λ multipliers on the minimum divergence problem $\Delta(\theta; P)$ in (17) differ from λ in criterions (13) and (14) by the factor η (see the discussion in Section 2.3). As our bootstrap methods are derived based on criterions (13) and (14), when implementing the bootstrap we rescale the multiplier λ solving (16) by the multiplier η on the constraint $\Delta(\theta; \hat{P}) \leq \delta$ in the outer optimization.[27] As η and λ are computed separately in the outer and inner optimizations, respectively, it is computationally most convenient to implement our bootstrap CSs with $\hat{\nu} = 0$. As discussed in Section 6.2, this choice is valid but possibly conservative. Despite this potentially conservative choice, the bootstrap CSs are not materially wider than the bootstrap CSs under the i.i.d. Gumbel assumption.

To construct the projection CSs, we form a 95% rectangular CS for P_{20} as described in Section 6.3. For each draw of $\hat{\theta}_\pi^*$ we compute the model-implied CCPs $\hat{P}_2^*$ under the i.i.d. Gumbel assumption. We construct t-statistics for each CCP by centering $\hat{P}_2^*$ at $\hat{P}_2$ and studentizing by its standard deviation across draws. For each draw we compute the

---

[27]This rescaling is also justified as follows. Let $\overline{b}_\delta(P) = \sup_{\theta \in \Theta : \Delta(\theta; P) \leq \delta} k(\theta)$ and note $\overline{\kappa}_\delta = \overline{b}_\delta(P_0)$ and $\hat{\overline{\kappa}}_\delta = \overline{b}_\delta(\hat{P})$. By similar arguments to Corollary 5 of Milgrom and Segal (2002), one may deduce that the directional derivative of $\overline{b}_\delta(P)$ at P_0 involves multiplying the directional derivative of $P \mapsto \Delta(\theta; P)$ at P_0 by the multiplier for $\Delta(\theta; P) \leq \delta$. The directional derivative of $P \mapsto \Delta(\theta; P)$ at P_0 may be shown to be

$$\lim_{n \to \infty} t_n^{-1} \left(\Delta(\theta; P_0 + t_n h_n) - \Delta(\theta; P_0) \right) = \sup_{\lambda_{12} \in \underline{\Lambda}(\theta; P_0)} -\lambda_{12}' h,$$

where $\underline{\Lambda}(\theta; P_0)$ is constructed analogously to $\underline{\Lambda}_\delta(\theta; P)$ in Section 6.2 using the set of multipliers that solve the minimum divergence problem (16).

Table 6: Computation times for the inner problem in the DDC application

			δ		
	0.01	0.1	1	10	100
Lower bound	0.124	0.144	0.164	0.285	0.265
Upper bound	0.101	0.119	0.142	0.266	1.039

Note: Times (in seconds) for solving the inner optimization problem at the parameter values at which $\hat{\underline{\kappa}}_\delta$ and $\hat{\bar{\kappa}}_\delta$ are attained. All computations are performed in Julia version 1.6.4 and Knitro 12.4.0 on a 2.7GHz MacBook Pro with 16GB memory.

maximum of the absolute value of the t-statistics. We then take the critical value $\hat{c}_{2,1-\alpha}$ to be the $1 - \alpha$ quantile of the maximum statistic across draws.

Computation times. Table 6 reports computation times for the inner optimization for evaluating the criterion functions $\hat{\underline{K}}_\delta(\theta; \hat{\gamma}, \hat{P})$ and $\hat{\bar{K}}_\delta(\theta; \hat{\gamma}, \hat{P})$ at the parameter values at which $\hat{\underline{\kappa}}_\delta$ and $\hat{\bar{\kappa}}_\delta$ are attained. The computation times correspond to solving the minimum divergence problem $\Delta(\theta; \hat{\gamma}, \hat{P})$ because k does not depend on u (cf. display (17)). The outer optimizations were typically solved in a few minutes in a 8-core environment with 64GB memory.

Sensitivity to ϕ. Bounds with χ^2-divergence were between 4% narrower and 1% wider than the bounds for hybrid divergence for all values of δ. Repeating the analysis with L^4-divergence, which is stronger than χ^2 and hybrid divergence, produced bounds that were 10-30% narrower than the hybrid divergence bounds up to $\delta = 1$ and at most 5% narrower than the hybrid divergence bounds for larger values of δ. As with the matching application, these results again show that the conclusions we draw from our analysis are not sensitive to the choice of ϕ function.

E   Proofs of Main Results

Throughout the proofs, we abbreviate upper-semicontinuous and upper-semicontinuity to u.s.c. and lower-semicontinuous and lower-semicontinuity to l.s.c.

E.1   Proofs for Section 2

Proof of Proposition 2.1. Immediate from Proposition G.1 in Appendix G.2 of Christensen and Connault (2022). ∎

Recall Condition S from Definition 2.1 and Condition S_{np} from Definition B.1.

Lemma E.1 *Suppose that Assumption Φ holds and μ and F_* are mutually absolutely continuous. Then Condition S holds at (θ, γ, P) if and only if Condition S_{np} holds at (θ, γ, P).*

Proof of Lemma E.1. In view of Hölder's inequality for Orlicz classes (see (A.1)), Assumption Φ implies $\mathcal{N}_\infty = \{F : D_\phi(F\|F_*) < \infty\} \subseteq \mathcal{F}_\theta$. Therefore,

$$\mathcal{G}_\infty := \{\mathbb{E}^F[g(U, \theta, \gamma)] : F \in \mathcal{N}_\infty\} \subseteq \{\mathbb{E}^F[g(U, \theta, \gamma)] : F \in \mathcal{F}_\theta\} =: \mathcal{G}_\theta.$$

By Corollary 6.6.2 of Rockafellar (1970), it suffices to show $\mathrm{ri}(\mathcal{G}_\infty) = \mathrm{ri}(\mathcal{G}_\theta)$. As $\mathrm{ri}(\mathcal{G}_\infty) \subseteq \mathcal{G}_\infty \subseteq \mathcal{G}_\theta$, it suffices to show $\mathcal{G}_\theta \subseteq \mathrm{cl}(\mathcal{G}_\infty)$ (Hiriart-Urruty and Lemaréchal, 2001, Remark 2.1.9). For any $x \in \mathcal{G}_\theta$, we have $x = \mathbb{E}^F[g(U, \theta, \gamma)]$ for some $F \in \mathcal{F}_\theta$. As $F \ll \mu$ and F_* and μ are mutually absolutely continuous, F has a density, say m, with respect to F_*. For each $n \geq 1$, let $m(u) \wedge n = \min\{m(u), n\}$ and define

$$m_n(u) = \frac{m(u) \wedge n}{\int (m(u) \wedge n) \, \mathrm{d}F_*(u)}.$$

Each F_n defined by $\mathrm{d}F_n = m_n \mathrm{d}F_*$ belongs to $\mathcal{N}_\infty$. It follows that $\mathbb{E}^{F_n}[g(U, \theta, \gamma)] \in \mathcal{G}_\infty$. By monotone convergence, we have $\mathbb{E}^{F_n}[g(U, \theta, \gamma)] \to x$. Therefore, $x \in \mathrm{cl}(\mathcal{G}_\infty)$. ∎

Proof of Theorem 2.1. We prove only the result for $\inf \mathcal{K}$; the result for $\sup \mathcal{K}$ follows similarly. Note

$$\inf \mathcal{K} = \inf_{\theta \in \Theta} \underline{K}_{np}(\theta; \gamma_0, P_0) = \inf_{\theta \in \Theta_I} \underline{K}_{np}(\theta; \gamma_0, P_0),$$

where the first equality is by definition and the second equality holds because, if $\theta \notin \Theta_I$, then there does not exist a distribution $F \in \mathcal{F}_\theta$ under which the moment conditions hold at (θ, γ_0, P_0) and consequently $\underline{K}_{np}(\theta; \gamma_0, P_0) = +\infty$. If $\theta \notin \Theta_I$, then there does not exist $F \in \mathcal{N}_\infty$ under which the moment conditions hold at (θ, γ_0, P_0) either because $\mathcal{N}_\infty \subseteq \mathcal{F}_\theta$ for all θ under Assumption Φ. Therefore, $\underline{K}_\infty(\theta; \gamma_0, P_0) = +\infty$ in that case too. We therefore have

$$\inf \mathcal{K}_\infty = \inf_{\theta \in \Theta_I} \underline{K}_\infty(\theta; \gamma_0, P_0).$$

In view of Lemma B.1, it suffices to show $\inf \mathcal{K} = \inf \mathcal{K}_\infty$. Note that $\inf \mathcal{K} \leq \inf \mathcal{K}_\infty$ holds by virtue of the inclusion $\mathcal{N}_\infty \subseteq \mathcal{F}_\theta$ for all θ. For the reverse inequality, choose any $\epsilon > 0$. By S-regularity of Θ_I, there exists $\underline{\theta} \in \Theta_I$ for which Condition S holds at $(\underline{\theta}, \gamma_0, P_0)$ and for which $\underline{K}_{np}(\underline{\theta}; \gamma_0, P_0) \leq \inf \mathcal{K} + \epsilon$. As Condition S holds at $(\underline{\theta}, \gamma_0, P_0)$ and $\mu \ll F_* \ll \mu$, Lemma E.1 implies that Condition S_{np} must also hold at $(\underline{\theta}, \gamma_0, P_0)$. Moreover, the μ-essential infimum and F_*-essential infimum of any function are equal because $\mu \ll$

$F_* \ll \mu$. Therefore by Lemmas B.2 and B.3, we have $\underline{K}_\infty(\underline{\theta}; \gamma_0, P_0) = \underline{K}_{np}(\underline{\theta}; \gamma_0, P_0)$. It follows by definition of $\inf \mathcal{K}_\infty$ that $\inf \mathcal{K}_\infty \leq \underline{K}_\infty(\underline{\theta}; \gamma_0, P_0) = \underline{K}_{np}(\underline{\theta}; \gamma_0, P_0) \leq \inf \mathcal{K} + \epsilon$. Therefore, $\inf \mathcal{K}_\infty \leq \inf \mathcal{K}$. ∎

E.2 Proofs for Section 3

Proof of Proposition 3.1. We prove the result for $\underline{K}_\delta$; the proof for $\overline{K}_\delta$ follows similarly. Consider

$$v^A := \inf_{\theta \in \Theta, F \in \mathcal{N}_\delta} \mathbb{E}^F[k(U, \theta, \gamma)] \quad \text{subject to (1) holding at } (\theta, \gamma, P), \qquad \text{(Program A)}$$

$$v^B := \inf_{\theta \in \Theta} \mathbb{E}^{\underline{F}_{\delta,\theta}}[k(U, \theta, \gamma)] \quad \text{subject to } \mathbb{E}^{\underline{F}_{\delta,\theta}}[g_{4e}(U, \theta, \gamma)] = 0, \qquad \text{(Program B)}$$

where $\underline{F}_{\delta,\theta}$ solves

$$\inf_{F \in \mathcal{N}_\delta} \mathbb{E}^F[k(U, \theta, \gamma)] \quad \text{subject to (19) holding at } (\theta, \gamma, P),$$

and $v^B = +\infty$ if there is no solution to this problem. Program A is the approach described in Section 2 whereas Program B is equivalent to our MPEC implementation.

The inequality $v^A \leq v^B$ is trivial if $v^B = +\infty$. If v^B is finite, for any $\varepsilon > 0$ there exists $\theta_\varepsilon^B \in \Theta$ for which $\mathbb{E}^{\underline{F}_{\delta,\theta_\varepsilon^B}}[k(U, \theta_\varepsilon^B, \gamma)] \leq v^B + \varepsilon$ and $\mathbb{E}^{\underline{F}_{\delta,\theta_\varepsilon^B}}[g_{4e}(U, \theta_\varepsilon^B, \gamma)] = 0$ where $\underline{F}_{\delta,\theta_\varepsilon^B}$ is well defined by Lemma G.2(ii) of Christensen and Connault (2022). As $(\theta_\varepsilon^B, \underline{F}_{\delta,\theta_\varepsilon^B})$ are feasible for Program A, we have $v^A \leq v^B + \varepsilon$. As ε is arbitrary, we have $v^A \leq v^B$ whenever $v^B > -\infty$.

A similar argument applies when $v^B = -\infty$: for any $n \in \mathbb{N}$ there exists $\theta_n^B \in \Theta$ for which $\mathbb{E}^{\underline{F}_{\delta,\theta_n^B}}[k(U, \theta_\varepsilon^B, \gamma)] \leq -n$ and $\mathbb{E}^{\underline{F}_{\delta,\theta_n^B}}[g_{4e}(U, \theta_\varepsilon^B, \gamma)] = 0$, where the distribution $\underline{F}_{\delta,\theta_n^B}$ is well defined by Lemma G.2(ii) of Christensen and Connault (2022). As $(\theta_n^B, \underline{F}_{\delta,\theta_n^B})$ are feasible for Program A, we have $v^A \leq -n$. Hence, $v^A = v^B = -\infty$.

Note $v^B \leq v^A$ holds trivially if $v^A = +\infty$. If v^A is finite, rewrite Program B as

$$\inf_{\kappa \in \mathbb{R}, \theta \in \Theta} \kappa \quad \text{subject to } \mathbb{E}^{\underline{F}_{\delta,\theta,\kappa}}[g_{4e}(U, \theta, \gamma)] = 0,$$

where $\underline{F}_{\delta,\theta,\kappa}$ solves the feasibility program

$$\inf_{F \in \mathcal{N}_\delta} 0 \quad \text{subject to (19) and } \mathbb{E}^F[k(U, \theta, \gamma)] = \kappa \text{ holding at } (\theta, \gamma, P). \qquad (49)$$

For any $\varepsilon > 0$ there exists $\theta_\varepsilon^A \in \Theta$ and $F_\varepsilon^A \in \mathcal{N}_\delta$ such that the constraints in Program A

are satisfied, i.e. $\mathbb{E}^{F_\varepsilon^A}[g_1(U, \theta_\varepsilon^A, \gamma)] \le P_1, \ldots, \mathbb{E}^{F_\varepsilon^A}[g_4(U, \theta_\varepsilon^A, \gamma)] = 0$, and

$$\mathbb{E}^{F_\varepsilon^A}[k(U, \theta_\varepsilon^A, \gamma)] \le v^A + \varepsilon.$$

Then $\underline{F}_\varepsilon^A$ solves the feasibility program (49) with $\theta = \theta_\varepsilon^A$ and $\kappa = \kappa_\varepsilon^A := \mathbb{E}^{F_\varepsilon^A}[k(U, \theta_\varepsilon^A, \gamma)]$. Note that $\mathbb{E}^{F_\varepsilon^A}[g_{4e}(U, \theta_\varepsilon^A, \gamma)] = 0$ also holds by construction. Therefore, $(\kappa_\varepsilon^A, \theta_\varepsilon^A)$ are feasible for the augmented form of Program B. It follows that $v^B \le \kappa_\varepsilon^A \le v^A + \varepsilon$ holds for each $\varepsilon > 0$. As $\varepsilon > 0$ is arbitrary, we have $v^B \le v^A$ whenever $v^A > -\infty$.

A similar argument applies if $v^A = -\infty$: for any $n \in \mathbb{N}$, we may choose $\theta_n^A \in \Theta$ and $F_n^A \in \mathcal{N}_\delta$ such that the constraints in Program A are satisfied and $\mathbb{E}^{F_n^A}[k(U, \theta_n^A, \gamma)] \le -n$. It follows that $v^B \le -n$. Hence, $v^B = v^A = -\infty$. ∎

Proof of Proposition 3.2. We prove the result for $\underline{F}_{\delta,\theta}$, the result for $\overline{F}_{\delta,\theta}$ follows similarly. We drop dependence of k and g on (θ, γ) to simplify notation in what follows.

First, suppose k depends on u. The dual formulation is justified by Proposition 2.1. A dual solution $(\underline{\eta}, \underline{\zeta}, \underline{\lambda})$ exists by Proposition G.1(iii) of Christensen and Connault (2022).

Suppose $\underline{\eta} > 0$. We wish to show that the change of measure $\underline{m}_{\delta,\theta}(u) = \dot{\phi}^\star(-\underline{\eta}^{-1}(k(u) + \underline{\zeta} + \underline{\lambda}'g_s(u)))$ induces a distribution that solves the primal problem (20) at θ. Differentiability of the objective function in (η, ζ, λ) is guaranteed by Assumption Φ. Also note that Assumption Φ(i) ensures $\dot{\phi}^\star \ge 0$. The first-order condition (FOC) for $\underline{\zeta}$ is

$$0 = \mathbb{E}^{F_*}\left[\dot{\phi}^\star(-\underline{\eta}^{-1}(k(U) + \underline{\zeta} + \underline{\lambda}'g_s(U)))\right] - 1$$

which implies $\mathbb{E}^{F_*}[\underline{m}_{\delta,\theta}] = 1$ and hence that $\underline{F}_{\delta,\theta}$ is a probability measure. The FOC for $\underline{\lambda}$ is

$$\begin{aligned}
0 &\ge \mathbb{E}^{F_*}\left[\dot{\phi}^\star(-\underline{\eta}^{-1}(k(U) + \underline{\zeta} + \underline{\lambda}'g_s(U)))g_1(U)\right] - P_1, \\
0 &= \mathbb{E}^{F_*}\left[\dot{\phi}^\star(-\underline{\eta}^{-1}(k(U) + \underline{\zeta} + \underline{\lambda}'g_s(U)))g_2(U)\right] - P_2, \\
0 &\ge \mathbb{E}^{F_*}\left[\dot{\phi}^\star(-\underline{\eta}^{-1}(k(U) + \underline{\zeta} + \underline{\lambda}'g_s(U)))g_3(U)\right], \\
0 &= \mathbb{E}^{F_*}\left[\dot{\phi}^\star(-\underline{\eta}^{-1}(k(U) + \underline{\zeta} + \underline{\lambda}'g_s(U)))g_{4s}(U)\right],
\end{aligned}$$

hence (1a)–(1c) and $\mathbb{E}^F[g_{4s}(U, \theta, \gamma)] = 0$ hold at (θ, γ, P) under $\underline{F}_{\delta,\theta}$. The FOC for $\underline{\eta} > 0$

17

is

$$0 = \mathbb{E}^{F_*}\left[\dot{\phi}^\star(-\underline{\eta}^{-1}(k(U) + \underline{\zeta} + \underline{\lambda}'g_s(U)))(-\underline{\eta}^{-1}(k(U) + \underline{\zeta} + \underline{\lambda}'g_s(U)))\right]$$
$$- \mathbb{E}^{F_*}\left[\phi^\star(-\underline{\eta}^{-1}(k(U) + \underline{\zeta} + \underline{\lambda}'g_s(U)))\right] - \delta\,.$$

By Assumption Φ(i), we may write the convex conjugate $\phi^{\star\star}$ of $\phi^\star$ using its Legendre transform:

$$\phi^{\star\star}(x^\star) = x^\star(\dot{\phi}^\star)^{-1}(x^\star) - \phi^\star((\dot{\phi}^\star)^{-1}(x^\star))$$

for any $x^\star$ in the range of $\dot{\phi}^\star$ (Rockafellar, 1970, Theorem 26.4). Setting $x^\star = \dot{\phi}^\star(x)$ and noting that $\phi^{\star\star} = \phi$ holds by the Fenchel–Moreau theorem, we obtain

$$\phi(\dot{\phi}^\star(x)) = x\dot{\phi}^\star(x) - \phi^\star(x)\,.$$

It follows that we may rewrite the FOC for $\underline{\eta}$ as $\delta = \mathbb{E}^{F_*}\left[\phi(\underline{m}_{\delta,\theta}(U))\right]$ and so $\underline{F}_{\delta,\theta} \in \mathcal{N}_\delta$.

Now suppose $\underline{\eta} = 0$. Here we wish to show that $\underline{m}_{\delta,\theta}(u) = \mathbb{1}\{u \in \underline{A}_{\delta,\theta}\}/F_*(\underline{A}_{\delta,\theta})$ induces a distribution that solves the primal problem (20) at θ. As the neighborhood constraint $F \in \mathcal{N}_\delta$ is not binding, the value of the objective must be the same as the optimal value when $\delta = \infty$. In view of Lemma B.2, the value is $F_*\text{-ess}\inf(k(\cdot) + \underline{\lambda}'g_s(\cdot)) - \underline{\lambda}'_{12}P$. We can write problem (22) as a nested optimization:

$$\sup_{\lambda \in \Lambda_s}\left(\sup_{\eta > 0, \zeta \in \mathbb{R}} -\eta\mathbb{E}^{F_*}\left[\phi^\star\left(\frac{k(U) + \zeta + \lambda'g_s(U)}{-\eta}\right)\right] - \eta\delta - \zeta - \lambda'_{12}P\right).$$

At $\lambda = \underline{\lambda}$, the inner problem is the dual of $\inf_{F \in \mathcal{N}_\delta} \mathbb{E}^F[k(U) + \underline{\lambda}'g_s(U) - \underline{\lambda}'_{12}P]$. As $\underline{\eta} = 0$, the constraint $F \in \mathcal{N}_\delta$ is not binding and so the minimizing distribution must be supported on $\underline{A}_{\delta,\theta}$. Finally, by convexity of ϕ, the distribution induced by $\underline{m}_{\delta,\theta}$ minimizes $D_\phi(\,\cdot\,\|F_*)$ among all distributions with support $\underline{A}_{\delta,\theta}$.

Now suppose k does not depend on u. By Proposition G.2 of Christensen and Connault (2022), the primal and dual values of (15) are equal and a dual solution exists. By similar arguments to above, $\mathbb{E}^{F_*}[\underline{m}_{\delta,\theta}(U)] = 1$, and (1a)–(1c) and $\mathbb{E}^F[g_{4s}(U, \theta, \gamma)] = 0$ hold at (θ, γ, P) under $\underline{F}_{\delta,\theta}$. Finally, as there exists $F \in \mathcal{N}_\delta$ under which the moment conditions (1a)–(1c) and $\mathbb{E}^F[g_{4s}(U, \theta, \gamma)] = 0$ hold at (θ, γ, P), we must have $D(\underline{F}_{\delta,\theta}\|F_*) \leq D(F\|F_*) \leq \delta$, as required. $\blacksquare$

E.3 Proofs for Section 4

Proof of Proposition 4.1. As $\phi_1(x) \leq \bar{a}\phi_2(x)$ for all $x > 0$, we have $D_{\phi_1}(F\|F_*) \leq \bar{a}D_{\phi_2}(F\|F_*)$. Hence, $\mathcal{N}_{\delta,2} \subseteq \mathcal{N}_{\bar{a}\delta,1}$ for each $\delta > 0$. The result follows from this inclusion, noting that $\underline{\kappa}_{\bar{a}\delta,1}$ and $\underline{\kappa}_{\bar{a}\delta,1}$ are both finite because Assumption Φ holds for ϕ_1. ∎

E.4 Proofs for Section 6

We first present some preliminary lemmas.

Lemma E.2 *Suppose that Assumptions Φ and $M(i),(v)$ hold. Let $\{(F_n, \theta_n, \gamma_n, P_n)\} \subseteq \mathcal{N}_\delta \times \Theta \times \Gamma \times \mathcal{P}$ with $(\gamma_n, P_n) \to (\tilde{\gamma}, \tilde{P}) \in \Gamma \times \mathcal{P}$ and with (1) holding under F_n at $(\theta_n, \gamma_n, P_n)$. Then: there exists a convergent subsequence $(F_{n_l}, \theta_{n_l}, \gamma_{n_l}, P_{n_l}) \to (\tilde{F}, \tilde{\theta}, \tilde{\gamma}, \tilde{P}) \in \mathcal{N}_\delta \times \Theta \times \Gamma \times \mathcal{P}$ along which $\lim_{l\to\infty} \mathbb{E}^{F_{n_l}}[k(U, \theta_{n_l}, \gamma_{n_l})] = \mathbb{E}^{\tilde{F}}[k(U, \tilde{\theta}, \tilde{\gamma})]$ and similarly for each entry of $g_1, \ldots, g_4$, and (1) holds under $\tilde{F}$ at $(\tilde{\theta}, \tilde{\gamma}, \tilde{P})$.*

Proof of Lemma E.2. Let $m_n = \frac{\mathrm{d}F_n}{\mathrm{d}F_*}$. By Assumption $M(v)$, $\{\theta_n\}$ has a convergent subsequence $\{\theta_{n_l}\}$. As $\{m_{n_l}\}$ is $\|\cdot\|_\phi$-norm bounded (Lemma F.1(ii) of Christensen and Connault (2022)), taking a further subsequence if necessary we may assume $\{m_{n_l}\}$ is $\mathcal{E}$-weakly convergent to $\tilde{m} \in \mathcal{L}$ (see Appendix F of Christensen and Connault (2022)). By the triangle inequality, the Hölder inequality (A.1), $\mathcal{E}$-weak convergence, and Assumption $M(i)$, we have

$$\left| \mathbb{E}^{F_{n_l}}[m_{n_l}(U)k(U, \theta_{n_l}, \gamma_{n_l})] - \mathbb{E}^{F_*}[\tilde{m}(U)k(U, \tilde{\theta}, \tilde{\gamma})] \right|$$
$$\leq |\mathbb{E}^{F_*}[(m_{n_l}(U) - \tilde{m}(U))k(U, \tilde{\theta}, \tilde{\gamma})]| + \|m_{n_l}\|_\phi \|k(\cdot, \theta_{n_l}, \gamma_{n_l}) - k(\cdot, \tilde{\theta}, \tilde{\gamma})\|_\psi \to 0.$$

It follows by similar arguments that

$$\mathbb{E}^{F_*}[\tilde{m}(U)] = 1, \qquad \mathbb{E}^{F_*}[\tilde{m}(U)g_1(U, \tilde{\theta}, \tilde{\gamma})] \leq \tilde{P}_1, \qquad \mathbb{E}^{F_*}[\tilde{m}(U)g_2(U, \tilde{\theta}, \tilde{\gamma})] = \tilde{P}_2,$$
$$\mathbb{E}^{F_*}[\tilde{m}(U)g_3(U, \tilde{\theta}, \tilde{\gamma})] \leq 0, \qquad \mathbb{E}^{F_*}[\tilde{m}(U)g_4(U, \tilde{\theta}, \tilde{\gamma})] = 0.$$

Finally, by Lemma F.1(i) of Christensen and Connault (2022), we have the inequality $\delta \geq \liminf_{l\to\infty} \mathbb{E}^{F_*}[\phi(m_{n_l}(U))] \geq \mathbb{E}^{F_*}[\phi(\tilde{m}(U))]$. ∎

Lemma E.3 *Suppose that Assumptions Φ and $M(i),(iii)-(v)$ hold. Then $\underline{\kappa}_\delta$ and $\overline{\kappa}_\delta$ are finite, and*

$$\underline{\kappa}_\delta = \inf_{\theta \in \Theta_\delta(\gamma_0, P_0)} \underline{K}_\delta(\theta; \gamma_0, P_0), \quad \overline{\kappa}_\delta = \sup_{\theta \in \Theta_\delta(\gamma_0, P_0)} \overline{K}_\delta(\theta; \gamma_0, P_0).$$

Proof of Lemma E.3. We prove the result only for $\underline{\kappa}_\delta$; the result for $\overline{\kappa}_\delta$ follows similarly.

Finiteness of $\underline{\kappa}_\delta$ follows by Assumptions Φ and M(i)(v) and the Hölder inequality (A.1). To simplify notation, we suppress dependence of $\Theta_\delta(\gamma_0, P_0)$ on (γ_0, P_0) in what follows. Suppose there is $\underline{\theta} \notin \Theta_\delta$ with $\underline{K}_\delta(\underline{\theta}; \gamma_0, P_0) < \inf_{\theta \in \Theta_\delta} \underline{K}_\delta(\theta; \gamma_0, P_0)$. Then there must exist $F_{\underline{\theta}} \in \mathcal{N}_\delta$ satisfying (1) at $(\underline{\theta}, \gamma_0, P_0)$. As $\Delta(\underline{\theta}; \gamma_0, P_0) = \delta$, it follows by convexity of ϕ that $F_{\underline{\theta}}$ must be unique. Therefore

$$\mathbb{E}^{F_{\underline{\theta}}}[k(U, \underline{\theta}, \gamma_0)] = \underline{K}_\delta(\underline{\theta}; \gamma_0, P_0) < \inf_{\theta \in \Theta_\delta} \underline{K}_\delta(\theta; \gamma_0, P_0) \le \inf_{\theta \in \Theta_\delta} \mathbb{E}^{F_\theta}[k(U, \theta_0, \gamma_0)], \qquad (50)$$

where, for each $\theta \in \Theta_\delta$, the distribution F_θ solves $\inf_F D_\phi(F \| F_*)$ subject to (1). Existence of F_θ follows by similar arguments to the proof of Lemma G.2(ii) of Christensen and Connault (2022); its uniqueness follows by strict convexity of ϕ.

Choose $\{\theta_n\} \subset \Theta_\delta$ with $\theta_n \to \underline{\theta}$ (we may choose such a sequence by Assumption M(iv)). By Lemma E.2, there is a subsequence $\{(\theta_{n_l}, F_{\theta_{n_l}}, \gamma_0, P_0)\}$ with $(\theta_{n_l}, F_{\theta_{n_l}}) \to (\underline{\theta}, \underline{F})$ for some $\underline{F} \in \mathcal{N}_\delta$ for which (1) holds under $\underline{F}$ at $(\underline{\theta}, \gamma_0, P_0)$. It follows by uniqueness of $F_{\underline{\theta}}$ that $\underline{F} = F_{\underline{\theta}}$. By Lemma E.2, we therefore have

$$\inf_{\theta \in \Theta_\delta} \mathbb{E}^{F_\theta}[k(U, \theta, \gamma_0)] \le \lim_{l \to \infty} \mathbb{E}^{F_{\theta_{n_l}}}[k(U, \theta_{n_l}, \gamma_0)] = \mathbb{E}^{F_{\underline{\theta}}}[k(U, \underline{\theta}, \gamma_0)],$$

which contradicts (50). ∎

Define

$$\underline{b}_\delta(\gamma, P) = \inf_{\theta \in \Theta_\delta(\gamma, P)} \underline{K}_\delta(\theta; \gamma, P), \quad \overline{b}_\delta(\gamma, P) = \inf_{\theta \in \Theta_\delta(\gamma, P)} \overline{K}_\delta(\theta; \gamma, P).$$

Lemma E.4 *Suppose that Assumptions Φ and M(i)–(v) hold. Then $\underline{b}_\delta(\gamma, P)$ and $\overline{b}_\delta(\gamma, P)$ are continuous at (γ_0, P_0).*

Proof of Lemma E.4. We prove the result only for $\underline{b}_\delta$; the result for $\overline{b}_\delta$ follows similarly.

Fix $\varepsilon > 0$. By Lemma E.3, we may choose $\theta_\varepsilon \in \Theta_\delta(\gamma_0, P_0)$ such that $\underline{K}_\delta(\theta_\varepsilon; \gamma_0, P_0) < \underline{b}_\delta(\gamma_0, P_0) + \varepsilon$. By Lemma G.8 of Christensen and Connault (2022) and Assumption M(ii) we have $\Delta(\theta_\varepsilon; \gamma, P) < \delta$ on a neighborhood N of (γ_0, P_0). Moreover, by Lemma G.9(i) of Christensen and Connault (2022) and Assumption M(i)–(iii) we have

$$\underline{K}_\delta(\theta_\varepsilon; \gamma, P) < \underline{K}_\delta(\theta_\varepsilon; \gamma_0, P_0) + \varepsilon$$

on a neighborhood N' of (γ_0, P_0). On $N \cap N'$ we therefore have

$$\underline{b}_\delta(\gamma, P) \le \underline{K}_\delta(\theta_\varepsilon; \gamma, P) < \underline{K}_\delta(\theta_\varepsilon; \gamma_0, P_0) + \varepsilon < \underline{b}_\delta(\gamma_0, P_0) + 2\varepsilon \,,$$

establishing u.s.c. of $\underline{b}_\delta(\gamma, P)$ at (γ_0, P_0).

To establish l.s.c., suppose there is $\varepsilon > 0$ and $(\gamma_n, P_n) \to (\gamma_0, P_0)$ along which

$$\underline{b}_\delta(\gamma_n, P_n) \le \underline{b}_\delta(\gamma_0, P_0) - 2\varepsilon \,. \tag{51}$$

Note $\Theta_\delta(\gamma_n, P_n)$ is nonempty for n sufficiently large by Lemma G.8 of Christensen and Connault (2022) and Assumption M(ii)(iii). For each n sufficiently large, choose $\theta_n \in \Theta_\delta(\gamma_n, P_n)$ and $F_n \in \mathcal{N}_\delta$ for which

$$\mathbb{E}^{F_n}[k(U, \theta_n, \gamma_n)] < \underline{b}_\delta(\gamma_n, P_n) + \varepsilon \,. \tag{52}$$

By Lemma E.2 there is a subsequence $(F_{n_l}, \theta_{n_l}, \gamma_{n_l}, P_{n_l}) \to (\underline{F}, \underline{\theta}, \gamma_0, P_0)$ for some $\underline{F} \in \mathcal{N}_\delta$ and $\underline{\theta} \in \Theta$, such that (1) holds under $\underline{F}$ at $(\underline{\theta}, \gamma_0, P_0)$, and for which

$$\lim_{l \to \infty} \mathbb{E}^{F_{n_l}}[k(U, \theta_{n_l}, \gamma_{n_l})] = \mathbb{E}^{\underline{F}}[k(U, \underline{\theta}, \gamma_0)] \ge \underline{K}_\delta(\underline{\theta}; \gamma_0, P_0) \,.$$

In view of (51) and (52) and Lemma E.3, this implies $\underline{K}_\delta(\underline{\theta}; \gamma_0, P_0) \le \underline{b}_\delta(\gamma_0, P_0) - \varepsilon = \underline{\kappa}_\delta - \varepsilon$, contradicting the definition of $\underline{\kappa}_\delta$. ∎

Proof of Theorem 6.1. Note that $\underline{\kappa}_\delta = \underline{b}_\delta(\gamma_0, P_0)$ and $\overline{\kappa}_\delta = \overline{b}_\delta(\gamma_0, P_0)$ by Lemma E.3 and $\hat{\underline{\kappa}}_\delta = \underline{b}_\delta(\hat\gamma, \hat P)$ and $\hat{\overline{\kappa}}_\delta = \overline{b}_\delta(\hat\gamma, \hat P)$ by definition. The result now follows by Lemma E.4 and Slutsky's theorem. ∎

Lemma E.5 *Suppose that Assumptions Φ and M(i),(ii) hold, Condition S' holds at (θ, γ, P), and $\Delta(\theta; \gamma, P) < \delta$. Then there is a neighborhood N of (θ, γ, P) such that Condition S' holds at $(\tilde\theta, \tilde\gamma, \tilde P)$ and $\Delta(\tilde\theta; \tilde\gamma, \tilde P) < \delta$ for all $(\tilde\theta, \tilde\gamma, \tilde P) \in N$.*

Proof of Lemma E.5. By Lemma G.7 of Christensen and Connault (2022), Condition S' holds at all $(\tilde\theta, \tilde\gamma, \tilde P)$ in a neighborhood N' of (θ, γ, P). Moreover, $\Delta(\tilde\theta; \tilde\gamma, \tilde P) < \delta$ holds at all $(\tilde\theta, \tilde\gamma, \tilde P)$ in a neighborhood N'' of (θ, γ, P) by Lemma G.8 of Christensen and Connault (2022). Set $N = N' \cap N''$. ∎

In the remainder of this subsection we drop dependence of all quantities on γ.

Proof of Theorem 6.2. We prove the result only for $\underline{b}_\delta$; the result for $\overline{b}_\delta$ follows similarly.

Step 1: We first show $\underline{\Theta}_\delta(P_0)$ is nonempty and compact. For nonemptiness, choose $\{\theta_n\}$ such that $\underline{K}_\delta(\theta_n; P_0) \downarrow \underline{\kappa}_\delta$. Let F_n solve the primal problem for θ_n. By Lemma E.2, there is a subsequence $(F_{n_l}, \theta_{n_l}) \to (\underline{F}, \underline{\theta})$ with $\underline{F} \in \mathcal{N}_\delta$ and $\underline{\theta} \in \Theta$ such that (1) holds under $\underline{F}$ at $(\underline{\theta}, P_0)$ and for which

$$\underline{\kappa}_\delta = \lim_{l \to \infty} \mathbb{E}^{F_{n_l}}[k(U, \theta_{n_l})] = \mathbb{E}^{\underline{F}}[k(U, \underline{\theta})].$$

Therefore, $\underline{\Theta}_\delta(P_0)$ is nonempty. We may deduce by similar arguments that $\underline{\Theta}_\delta(P_0)$ is closed. Compactness now follows by Assumption M(v).

Step 2: We now prove directional differentiability. Let $P_n = P_0 + t_n h_n$ with $t_n \downarrow 0$ and $h_n \to h$. Choose $\underline{\theta} \in \underline{\Theta}_\delta(P_0)$. By Lemma E.5 and Assumption M(iii)(vi), Condition S' holds at $(\underline{\theta}, P_n)$ and $\Delta(\underline{\theta}; P_n) < \delta$ for n sufficiently large, so by Proposition G.1(iv) of Christensen and Connault (2022) the set $\underline{\Lambda}_\delta(\underline{\theta}; P_n)$ is nonempty and compact for n sufficiently large. It now follows by definition of the objective (13) that

$$\underline{b}_\delta(P_n) - \underline{b}_\delta(P_0) \le \underline{K}_\delta(\underline{\theta}; P_n) - \underline{K}_\delta(\underline{\theta}; P_0) \le t_n \times -\underline{\lambda}'_{12} h_n,$$

for all $\underline{\lambda}_{12} \in \underline{\Lambda}_\delta(\underline{\theta}; P_n)$. Finally, by Lemma G.9(ii) of Christensen and Connault (2022) we obtain

$$\limsup_{n \to \infty} \frac{\underline{b}_\delta(P_n) - \underline{b}_\delta(P_0)}{t_n} \le \max_{\underline{\lambda}_{12} \in \underline{\Lambda}_\delta(\underline{\theta}; P_0)} -\underline{\lambda}'_{12} h.$$

Taking the infimum of both sides over $\underline{\theta} \in \underline{\Theta}_\delta$ yields

$$\limsup_{n \to \infty} \frac{\underline{b}_\delta(P_n) - \underline{b}_\delta(P_0)}{t_n} \le \inf_{\theta \in \underline{\Theta}_\delta} \max_{\underline{\lambda}_{12} \in \underline{\Lambda}_\delta(\theta; P_0)} -\underline{\lambda}'_{12} h. \tag{53}$$

For the lower bound, choose $\theta_n \in \underline{\Theta}_\delta(P_n)$ with $\underline{K}_\delta(\theta_n; P_n) \le \underline{b}_\delta(P_n) + t_n^2$ for all n sufficiently large. Take a subsequence $\{\theta_{n_l}\}$. By Assumption M(v) (taking a further subsequence if necessary), we have $\theta_{n_l} \to \underline{\theta} \in \Theta$. By similar arguments to step 1, we may in fact deduce that $\underline{\theta} \in \underline{\Theta}_\delta$. Reasoning as above, for l sufficiently large we have

$$\underline{b}_\delta(P_{n_l}) - \underline{b}_\delta(P_0) \ge \underline{K}_\delta(\theta_{n_l}; P_{n_l}) - \underline{K}_\delta(\theta_{n_l}; P_0) - t_{n_l}^2 \ge t_{n_l} \times -\underline{\lambda}'_{12} h_{n_l} - t_{n_l}^2,$$

where the final inequality holds for any $\underline{\lambda}_{12} \in \underline{\Lambda}_\delta(\theta_{n_l}; P_0)$. By Assumption M(vii), we may choose $\underline{\lambda}_{12,n_l} \in \underline{\Lambda}_\delta(\theta_{n_l}; P_0)$ for which $-\underline{\lambda}'_{12,n_l} h \to \max_{\underline{\lambda}_{12} \in \underline{\Lambda}_\delta(\underline{\theta}; P_0)} -\underline{\lambda}'_{12} h$ as $l \to \infty$.

Therefore,

$$\liminf_{l\to\infty} \frac{\underline{b}_\delta(P_{n_l}) - \underline{b}_\delta(P_0)}{t_{n_l}} \geq \max_{\underline{\lambda}_{12}\in\underline{\Lambda}_\delta(\theta;P_0)} -\underline{\lambda}'_{12}h \geq \inf_{\theta\in\underline{\Theta}_\delta} \max_{\underline{\lambda}_{12}\in\underline{\Lambda}_\delta(\theta;P_0)} -\underline{\lambda}'_{12}h\,.$$

As the lower bound does not depend on the subsequence $\{\theta_{n_l}\}$, we have

$$\liminf_{n\to\infty} \frac{\underline{b}_\delta(P_n) - \underline{b}_\delta(P_0)}{t_n} \geq \inf_{\theta\in\underline{\Theta}_\delta} \max_{\underline{\lambda}_{12}\in\underline{\Lambda}_\delta(\theta;P_0)} -\underline{\lambda}'_{12}h\,, \tag{54}$$

proving directional differentiability. Finally, Assumption M(vii) and Lemma G.9(ii) of Christensen and Connault (2022) imply $\theta \mapsto \underline{\Lambda}_\delta(\theta; P_0)$ is continuous at each $\theta \in \underline{\Theta}_\delta$. The set $\underline{\Lambda}_\delta(\theta; P_0)$ is also compact for each $\theta \in \underline{\Theta}_\delta$ by Proposition G.1(iv) of Christensen and Connault (2022). It follows by the maximum theorem that the infima in (53) and (54) can be replaced by minima.

Step 3: In view of step 2, the asymptotic distribution follows by Theorem 2.1 of Shapiro (1991) and the fact that $\sqrt{n}(\hat{P} - P) \to_d N(0,\Sigma)$. ∎

Proof of Theorem 6.3. We verify the conditions of Theorem 3.2 of Fang and Santos (2019). Their Assumptions 1 and 2 hold by Theorem 6.2 and because $\sqrt{n}(\hat{P} - P_0) \to_d N(0,\Sigma)$ with Σ finite, respectively. Their Assumption 3 is assumed directly. Finally, Lemma G.11 of Christensen and Connault (2022) shows that $\widehat{d\underline{b}}_{\delta,P_0}$ and $\widehat{d\overline{b}}_{\delta,P_0}$ satisfy the sufficient conditions for Assumption 4 of Fang and Santos (2019), which is presented in their Remark 3.4. This proves consistency. Coverage of $CS_{\delta,L}^{1-\alpha}$ and $CS_{\delta,U}^{1-\alpha}$ follows by continuity of the distribution functions. Coverage of $CS_\delta^{1-\alpha}$ follows by the Bonferroni inequality. ∎

Proof of Theorem 6.4. We prove the result only for $CS_\delta^{1-\alpha}$; the result for the other CSs follow similarly. Say that $P_0 \in CS_{P_0}^{1-\alpha}$ if $P_{10} \leq \hat{P}_{1,U}^{1-\alpha}$ and $P_{20} \in [\hat{P}_{2,L}^{1-\alpha}, \hat{P}_{2,U}^{1-\alpha}]$ both hold. By Lemma E.3, for each $\varepsilon > 0$ we may choose $\underline{\theta}_\varepsilon, \overline{\theta}_\varepsilon \in \Theta_\delta(P_0)$ such that $\underline{K}_\delta(\underline{\theta}_\varepsilon; P_0) < \underline{\kappa}_\delta + \varepsilon$ and $\overline{K}_\delta(\overline{\theta}_\varepsilon; P_0) > \overline{\kappa}_\delta - \varepsilon$. Let $\underline{F}_{\underline{\theta}_\varepsilon}$ and $\overline{F}_{\overline{\theta}_\varepsilon}$ solve problem (15) at $(\underline{\theta}_\varepsilon; P_0)$ and $(\overline{\theta}_\varepsilon; P_0)$, respectively. Whenever $P_0 \in CS_{P_0}^{1-\alpha}$ holds, $\underline{F}_{\underline{\theta}_\varepsilon}$ and $\overline{F}_{\overline{\theta}_\varepsilon}$ must also satisfy the "relaxed" moment conditions used for computing $\hat{\underline{\kappa}}_{\delta,1-\alpha}$ and $\hat{\overline{\kappa}}_{\delta,1-\alpha}$, so it follows that $\Delta_{cs}(\underline{\theta}_\varepsilon; \hat{P}_{1-\alpha}) < \delta$ and $\Delta_{cs}(\overline{\theta}_\varepsilon; \hat{P}_{1-\alpha}) < \delta$. Moreover, as the primal solutions for $\underline{K}_\delta(\underline{\theta}_\varepsilon; P_0)$ and $\overline{K}_\delta(\overline{\theta}_\varepsilon; P_0)$ are feasible for the relaxed problem whenever $P_0 \in CS_{P_0}^{1-\alpha}$, we have

$$\hat{\underline{\kappa}}_{\delta,1-\alpha} \leq \underline{K}_{\delta,cs}(\underline{\theta}_\varepsilon; \hat{P}_{1-\alpha}) \leq \underline{K}_\delta(\underline{\theta}_\varepsilon; P_0) < \underline{\kappa}_\delta + \varepsilon\,,$$

and similarly $\hat{\overline{\kappa}}_{\delta,1-\alpha} > \overline{\kappa}_\delta - \varepsilon$. As ε is arbitrary, we have that $\underline{\kappa}_\delta \geq \hat{\underline{\kappa}}_{\delta,1-\alpha}$ and $\overline{\kappa}_\delta \leq \hat{\overline{\kappa}}_{\delta,1-\alpha}$

holds whenever $P_0 \in CS_{P_0}^{1-\alpha}$. The desired coverage now follows by (30). ∎

References

Andrews, I., M. Gentzkow, and J. M. Shapiro (2017). Measuring the sensitivity of parameter estimates to estimation moments. *The Quarterly Journal of Economics 132*(4), 1553–1592.

Andrews, I., M. Gentzkow, and J. M. Shapiro (2020). On the informativeness of descriptive statistics for structural estimates. *Econometrica* (6), 2231–2258.

Bonhomme, S. and M. Weidner (2021). Minimizing sensitivity to model misspecification. *Quantitative Economics*, forthcoming.

Christensen, T. and B. Connault (2022). Counterfactual sensitivity and robustness. arxiv preprint, arxiv:1904.00989 [econ.em].

Csiszár, I. and P. C. Shields (2004). Information theory and statistics: A tutorial. *Communications and Information Theory 1*(4), 417–528.

Fang, Z. and A. Santos (2019). Inference on directionally differentiable functions. *Review of Economic Studies 86*(1), 377–412.

Hiriart-Urruty, J.-B. and C. Lemaréchal (2001). *Fundamentals of Convex Analysis*. Springer.

Lehmann, E. L. and G. Casella (1998). *Theory of Point Estimation* (2 ed.). Springer-Verlag New York, Inc.

Milgrom, P. and I. Segal (2002). Envelope theorems for arbitrary choice sets. *Econometrica 70*(2), 583–601.

Rockafellar, R. T. (1970). *Convex Analysis*. Princeton University Press.

Shapiro, A. (1991). Asymptotic analysis of stochastic programs. *Annals of Operations Research 30*(1), 169–186.