

THEORY OF WEAK IDENTIFICATION IN SEMIPARAMETRIC MODELS

TETSUYA KAJI

University of Chicago Booth School of Business

We provide general formulation of weak identification in semiparametric models and an efficiency concept. Weak identification occurs when a parameter is weakly regular, that is, when it is locally homogeneous of degree zero. When this happens, consistent or equivariant estimation is shown to be impossible. We then show that there exists an underlying regular parameter that fully characterizes the weakly regular parameter. While this parameter is not unique, concepts of sufficiency and minimality help pin down a desirable one. If estimation of minimal sufficient underlying parameters is inefficient, it introduces noise in the corresponding estimation of weakly regular parameters, whence we can improve the estimators by local asymptotic Rao–Blackwellization. We call an estimator weakly efficient if it does not admit such improvement. New weakly efficient estimators are presented in linear IV and nonlinear regression models. Simulation of a linear IV model demonstrates how 2SLS and optimal IV estimators are improved.

KEYWORDS: Weak identification, semiparametric efficiency.

1. INTRODUCTION

WEAK IDENTIFICATION arises in a wide range of empirical settings. A leading example is the linear instrumental variables (IV) model in which the instruments and endogenous regressors are barely correlated. When this happens, classical asymptotic theory is known to yield poor approximations to the behavior of familiar statistics (Staiger and Stock (1997)). We encounter this problem in various other contexts: Stock and Wright (2000) analyzed weak identification in generalized method of moments (GMM) models; Andrews and Cheng (2012) in extremum estimation models; Andrews and Mikusheva (2016) in dynamic stochastic general equilibrium (DSGE) models. Many estimators of weakly identified parameters exhibit inconsistency and bias, and, as a consequence, standard inference procedures such as t - and Wald tests may have substantially distorted sizes (Dufour (1997), Hirano and Porter (2015)). Following these problems, a vast amount of theoretical work has been published.

The theoretical literature on weak identification is largely confined to specific estimation and inference procedures in specific models. Many fundamental questions—such as what is the common cause of weak identification, what is a general guideline to find well-behaved statistics, and what is semiparametric efficiency in the presence of weak identification—are left unanswered. Such exploration is essential, however, not only to facilitate unified understanding of the phenomenon but to measure performance of different procedures and develop systematic methods for estimation and inference.

This paper studies weak identification from the perspective of semiparametric theory (Bickel et al. (1993), Van der Vaart (1998, Chapter 25)). This literature views parameters as functions defined on the probability manifold and relates their asymptotic properties to

Tetsuya Kaji: tkaji@chicagobooth.edu

I thank Anna Mikusheva, Isaiah Andrews, Victor Chernozhukov, Whitney Newey, Jerry Hausman, Hidehiko Ichimura, Kengo Kato, Rachael Meager, Peter Hull, and anonymous referees for helpful comments and suggestions. This work is supported by the Richard N. Rosett Faculty Fellowship and the Liew Family Faculty Fellowship at the University of Chicago Booth School of Business. A replication file is posted (Kaji (2021)).

the functions' local behaviors in response to local perturbations of the probability. While strongly identified parameters often translate to differentiable functions, weakly identified parameters emerge as functions that are discontinuous at the probability to which we asymptote. We call such functions *weakly regular* in reference to differentiable functions called *regular*. As an immediate consequence of this discontinuity, we derive—without reference to a specific estimation or inference procedure—that there exists neither a consistent estimator, a consistent test, nor an equivariant (hence pivotal) estimator when the parameter is weakly regular. The local approximations of weakly regular parameters are homogeneous of degree zero and essentially nonlinear, becoming the root cause of non-Gaussian nonpivotal asymptotic distributions witnessed throughout the literature (Staiger and Stock (1997), Stock and Wright (2000), Andrews and Cheng (2012)). To circumvent the problem of nonlinearity, we explore weak regularity from the standpoint of *regular* parameters.

We show that every weakly regular parameter can be represented as a nonlinear function of the local parameter of some underlying regular parameter. Finding such a parameter lets us reformulate the model in a way that it consists only of regular parameters, providing a tractable foundation on which to discuss estimation and inference. This conforms with the repeated observation in the literature that reduction to regular parameters (usually referred to as “reduced-form parameters”) substantially simplifies the problems; we generalize this observation to arbitrary semiparametric models and show that there exists an underlying regular parameter for every weakly regular parameter. However, underlying regular parameters are not unique, and statistical analyses based on different underlying parameters may yield different performances. This gives rise to the need for criteria to choose an underlying parameter.

We introduce two desirable properties of underlying parameters. In semiparametric models, the space of probability distributions can be much bigger than the space of the parameter of interest. Consequently, there are many directions of perturbations that do not matter to the parameter of interest. Intuitively, a good underlying parameter would be sensitive to all perturbations that matter *and* not sensitive to any perturbations that are irrelevant. The first property, *sufficiency*, ensures that all perturbations that affect identification of the weakly regular parameter affect the underlying parameter; for example, reduced-form coefficients in linear IV that miss an instrument are not sufficient. The second property, *minimality*, guarantees that the underlying parameter does not include a nuisance parameter. In short, a best underlying regular parameter is minimal and sufficient. We show existence of minimal sufficient underlying parameters and present examples.

With these concepts, we define a notion of efficiency for weakly regular parameters. Efficient estimation under weak identification has received little treatment as non-Gaussianity and nonpivotality of the asymptotic distributions render the classical efficiency concepts inapplicable in their direct forms. Our formulation helps decompose estimation of weakly regular parameters into estimation of the minimal sufficient underlying regular parameters and their transformation. Being regular, the underlying parameters are susceptible to the classical convolution theorem. Moreover, if the estimators of the underlying parameters contain noise, their transformations, too, would suffer from noise. Such noise can then be eliminated by taking expectation with respect to it. Conceptually, this corresponds to applying the Rao–Blackwell theorem to the local asymptotic representations of the estimators. The resulting conditional expectation estimators are more concentrated around the same means. We name this procedure *local asymptotic Rao–Blackwellization (LAR)*. If such improvement is impossible, we call the estimators *weakly*

efficient. We put the qualifier “weakly” as weakly efficient estimators are not unique. We can interpret weak efficiency as a generalization of classical efficiency to accommodate non-differentiable transformations.

We apply our results to linear IV and nonlinear regression models and present weakly efficient estimators. In linear IV, the two-stage least squares (2SLS) and even the optimal IV estimators are shown to be inefficient in the presence of heteroskedasticity and, under the availability of an efficient estimator of the reduced-form coefficients, admit transformations into weakly efficient estimators. In nonlinear regression, a simple least-squares estimator is shown to be inefficient under heteroskedasticity, and we obtain a weakly efficient estimator when the heteroscedastic structure can be estimated. In nonlinear GMM, the possibility of improvement depends on specificity of the model. Simulation shows how weakly efficient estimators behave under weak and strong identification asymptotics in a linear IV model.

There is a large body of literature that studies the optimality of statistical procedures under weak identification. To name a few, Müller and Wang (2019) studied estimation that minimizes the weighted average risk when the asymptotic distribution of the statistics is known; Andrews, Moreira, and Stock (2006) developed optimal conditional likelihood ratio tests in linear IV models with normal homoscedastic errors; Müller (2011) studied efficient inference under a weak convergence assumption.

The rest of the paper is organized as follows. Section 2 defines weak identification in semiparametric models, gives impossibility results, and introduces the notion of underlying regular parameters. Section 3 introduces sufficiency and minimality of underlying regular parameters. Section 4 derives LAR for the estimation of weakly regular parameters, whence we define weak efficiency. Section 5 discusses application of LAR to heteroscedastic linear IV models and provides simulation results. Section 6 concludes. The Appendix contains proofs and the local-to-singularity linear IV model.

2. WEAK IDENTIFICATION IN SEMIPARAMETRIC MODELS

Weak identification is a problem that arises relative to the sample size. As such, it is modeled by one of two ways: (1) start with a primitive (global) model and consider a drifting sequence, “weak identification embedding,” toward identification failure (e.g., Staiger and Stock (1997)); (2) skip the primitive model and start with a finite-sample situation equivalent to a local expansion around identification failure (e.g., Dufour (1997)). Both approaches are inherently local, although distinction may seem ambiguous in linear models where global and local settings coincide. To establish a framework for efficiency, it is essential to understand the type of local expansions that can emerge from a primitive model. Therefore, we develop a general theory for (1).

Suppose we observe i.i.d. random variables X_1, \dots, X_n from the sample space $(\mathcal{X}, \mathcal{A})$. The set of possible distributions of X_i is denoted by \mathcal{P} and called the *model*. To obtain fruitful asymptotics around a distribution $P \in \mathcal{P}$, we consider a *path* of distributions $Q_t \in \mathcal{P}$ indexed by a real number $t \in (0, 1]$ that is *differentiable in quadratic mean (DQM)* at P , that is, there exists a measurable function $g : \mathcal{X} \rightarrow \mathbb{R}$ such that

$$\int_{\mathcal{X}} \left[\frac{dQ_t^{1/2} - dP^{1/2}}{t} - \frac{1}{2}g dP^{1/2} \right]^2 \rightarrow 0 \quad \text{as } t \rightarrow 0,$$

where the integral is understood with respect to a σ -finite measure dominating P and Q_t , and dP and dQ_t are their Radon–Nikodym derivatives. This convergence is denoted by

$Q_t \rightarrow^{\text{DQM}} P$, and we call g the (*model*) *score* induced by the path $\{Q_t\}$.¹ The idea behind asymptotic approximation theory is that the path of “alternatives” $\{Q_t\}$ that approaches P at the same rate as the path of “samples” $\{\hat{P}_n\}$ is not deterministically distinguishable in the limit and hence yields an approximation that reflects finite-sample uncertainty; therefore, $t = 1/\sqrt{n}$ under local asymptotic normality (Van der Vaart (1998, Lemma 25.14)), and in a minor abuse of notation we denote $Q_{1/\sqrt{n}}$ by Q_n .

We often do not consider every possible path in \mathcal{P} ; let \mathcal{P}_P denote the set of paths we consider that tend to P in DQM. Since there is little chance of misunderstanding, we denote $\{Q_t\}$ simply by Q_t , for example, $Q_t \in \mathcal{P}_P$; therefore, Q_t can refer to the entire path $\{Q_t\}$ or an element Q_t for a specific t , depending on the context. The set $\dot{\mathcal{P}}_P$ of scores g induced by the paths in \mathcal{P}_P is called the *tangent set* at P . It is clear that $\dot{\mathcal{P}}_P$ is a subset of zero-mean functions in $L_2(P)$.² Depending on the structures of \mathcal{P} and \mathcal{P}_P , the tangent set might be a linear space, a cone,³ or a set without much structure; we assume that \mathcal{P} and \mathcal{P}_P are nice enough that the induced tangent set is linear. For this reason, we call the tangent set the *tangent space*. The tangent space can be considered the local approximation of the model by a linear vector space. Finally, a parameter $\psi : \mathcal{P} \rightarrow \mathbb{D}$ is defined as a map from the model \mathcal{P} to a Banach space \mathbb{D} .

If the parameter $\psi : \mathcal{P} \rightarrow \mathbb{D}$ is differentiable in a suitable sense, we may approximate the change in the parameter along a path by a linear map from the tangent space $\dot{\mathcal{P}}_P$ to \mathbb{D} . Any infinitesimal perturbation of distribution P then leads to a linear perturbation of the parameter ψ . Such a parameter is known to behave well and is said to be *regular*. The appropriate notion of differentiability is as follows.

DEFINITION—Regular Parameter: A parameter $\psi : \mathcal{P} \rightarrow \mathbb{D}$ is *regular* (or *differentiable*) at $P \in \mathcal{P}$ relative to \mathcal{P}_P if there exists a continuous linear map $\dot{\psi}_P : \dot{\mathcal{P}}_P \rightarrow \mathbb{D}$ such that

$$\frac{\psi(Q_t) - \psi(P)}{t} \rightarrow \dot{\psi}_P g \quad \text{for every } Q_t \in \mathcal{P}_P.$$

The derivative map $\dot{\psi}_P$ is called the *local parameter* of ψ . Its adjoint map $\dot{\psi}_P^* : \mathbb{D}^* \rightarrow \overline{\dot{\mathcal{P}}_P}$ is called the *efficient influence map* of ψ , where \mathbb{D}^* is the dual space of \mathbb{D} and $\overline{\dot{\mathcal{P}}_P}$ the completion of $\dot{\mathcal{P}}_P$.⁴

REMARK: In the classical context, the tangent set “represents” the set of paths, so regularity (differentiability) is often defined “relative to the tangent set.” In the context of weak identification, however, the corresponding tangent set does not represent the set of paths; therefore, we keep the original wording “relative to the set of paths” from Van der Vaart (1991b). The term “regular” is taken from Van der Vaart and Wellner (1996, Chapter 3.11).

¹Throughout the paper, dependence of g on $\{Q_t\}$ will be implied by the context.

²In this sense, $\dot{\mathcal{P}}$ is the set of *equivalence classes* of scores, to be precise.

³A subset X of a linear space is called a *cone* if $x \in X$ implies $ax \in X$ for every $a > 0$.

⁴If there is a function $\tilde{\psi}_P : \mathcal{X} \rightarrow \mathbb{D}$ such that $\dot{\psi}_P^* \delta^* = \delta^* \tilde{\psi}_P$ for every $\delta^* \in \mathbb{D}^*$, it is called the *efficient influence function* (Bickel et al. (1993, Section 5.2)).

2.1. *Weakly Regular Parameters*

Now we define a weakly identified parameter. Let \mathcal{P}_β be a subset of \mathcal{P} on which a parameter β is uniquely defined.⁵ As the problem of weak identification arises when the population distribution is close to a point of identification failure, we model the situation by a path that takes values in \mathcal{P}_β and approaches a point outside of \mathcal{P}_β . However, not all such sequences are appropriate to consider. If the path approaches $\mathcal{P} \setminus \mathcal{P}_\beta$ too rapidly, β may not be identified in the first-order local expansion (tangent space) of \mathcal{P} . To avoid this, we focus on scores that are associated with unique limiting values of β .

DEFINITION—Pertinent Tangent Cone: The tangent set $\dot{\mathcal{P}}_{P,\beta} \subset \dot{\mathcal{P}}_P$ *pertinent to the sub-model* \mathcal{P}_β at $P \in \mathcal{P}$, possibly $P \in \mathcal{P} \setminus \mathcal{P}_\beta$, is the set of scores $g \in \dot{\mathcal{P}}_P$ such that there exists a path in \mathcal{S}_P that takes values in \mathcal{P}_β and induces g and every such path shares the same limit of $\beta(Q_t)$. Define $\mathcal{S}_{P,\beta}$ to be the set of paths in \mathcal{S}_P that take values in \mathcal{P}_β and induce scores in $\dot{\mathcal{P}}_{P,\beta}$.

Consequently, this paper does not cover faster-than- \sqrt{n} weak identification. From the observation that P is not in \mathcal{P}_β , we see that $\dot{\mathcal{P}}_{P,\beta}$ is only a cone.

LEMMA 1: $\dot{\mathcal{P}}_{P,\beta}$ and $\dot{\mathcal{P}}_P \setminus \dot{\mathcal{P}}_{P,\beta}$ are cones.

REMARK: In classical asymptotic theory, the limit distribution P is often regarded as the “null hypothesis” and the path Q_t as a drifting sequence of “alternatives.” When it comes to weak identification, both the null and alternatives reside as paths in $\mathcal{S}_{P,\beta}$; P is merely a point of reference for identification failure.

If the set of paths \mathcal{S}_P is much richer than $\mathcal{S}_{P,\beta}$ in a way that $\text{Span } \dot{\mathcal{P}}_{P,\beta}$ is a strict subset of $\dot{\mathcal{P}}_P$, then there exists a superfluously rich side of the model on which β is not even defined. Since it is meaningless to consider such parts of the model when one’s focus is on the parameter β , we assume innocuously that $\overline{\text{Span } \dot{\mathcal{P}}_{P,\beta}} = \overline{\dot{\mathcal{P}}_P}$.⁶

Now we define the weakly identified parameter under the name *weakly regular parameter*. We henceforth shun the use of the qualifier “weakly identified” since weak identification in the literature may not always exclude cases of, in fact, *no* identification. In this paper, we require that weakly regular parameters are identified at every fixed n in that there exists a unique value of the parameter for any given Q_n along a path. Moreover, we require that the parameters remain identified in the limit in the sense that there exists a unique limiting value of the parameter for each score g in $\dot{\mathcal{P}}_{P,\beta}$. Let \mathbb{B} be another Banach space on which a weakly regular parameter will be defined.

DEFINITION—Weakly Regular Parameter: A parameter $\beta : \mathcal{P}_\beta \rightarrow \mathbb{B}$ is *weakly regular* at $P \in \mathcal{P}$, possibly $P \in \mathcal{P} \setminus \mathcal{P}_\beta$, relative to $\mathcal{S}_{P,\beta}$ if there exists a map $\beta_P : \dot{\mathcal{P}}_{P,\beta} \rightarrow \mathbb{B}$ that is continuous on $\dot{\mathcal{P}}_{P,\beta}$ (not necessarily on $\dot{\mathcal{P}}_P$) and homogeneous of degree zero such that

$$\beta(Q_t) \longrightarrow \beta_P(g) \quad \text{for every } Q_t \in \mathcal{S}_{P,\beta}.$$

⁵ $\mathcal{P} \setminus \mathcal{P}_\beta$ may contain distributions that we simply deem inconceivable as well as distributions that do not identify β .

⁶Later on, we define the underlying regular parameter on the whole of \mathcal{P} , so it is actually harmful to require that the parameter be regular on the unconsidered realm of the model.

The definition says that the value to which a weakly regular parameter converges changes as we consider different paths. Moreover, this dependence is homogeneous of degree zero, hence essentially nonlinear and discontinuous at $g = 0$. This makes consistent estimation impossible and asymptotic distribution nonstandard (Section 2.2).

REMARK: Being a continuous map, a regular parameter is trivially weakly regular; that is, if $\psi : \mathcal{P} \rightarrow \mathbb{D}$ is regular, then $\psi(Q_t) \rightarrow \psi_P(g)$ where $\psi_P(g) \equiv \psi(P)$. Also, if β is a nontrivial weakly regular parameter, that is, β_P is nonconstant, then β_P cannot be linear since a linear function that is homogeneous of degree zero must be identically zero.

REMARK: Homogeneity of β_P is a natural consequence of dependence on g . Since $\beta(Q_{kt})$ for fixed $k > 0$ converges to the same limit as $\beta(Q_t)$, we have $\beta_P(kg) = \beta_P(g)$.

Now we introduce examples and show how they satisfy our definition.

EXAMPLE 1—Linear IV: Consider the IV regression model:

$$\begin{cases} y_i = x'_i\beta + \varepsilon_i = z'_i\pi\beta + u_i, & \mathbb{E}[\varepsilon_i | z_i] = 0, & \mathbb{E}[u_i | z_i] = 0, \\ x'_i = z'_i\pi + v'_i, & \mathbb{E}[v_i | z_i] = 0, \end{cases}$$

where y_i, ε_i, u_i are scalars, x_i, β, v_i are $d \times 1$ vectors, z_i is a $k \times 1$ vector, π is a $k \times d$ full column rank matrix, and $k \geq d$. We show that β is weakly regular under standard assumptions and the local-to-zero asymptotics of Staiger and Stock (1997). The local-to-singularity asymptotics is discussed in Appendix B.

Let \mathcal{P}_{uvz} be the set of probability distributions P_{uvz} on (u, v', z) with second moments such that $\mathbb{E}[u | z] = 0, \mathbb{E}[v | z] = 0, \mathbb{E}[zz']$ and $\mathbb{E}[vv' | z]$ are invertible, P_{uvz} dominated by the Lebesgue measure, and dP_{uvz} differentiable almost everywhere in (u, v') .⁷ The model \mathcal{P} is the set of distributions P on observables (x, y, z) such that

$$dP(x, y, z) = dP_{uvz}(y - z'\gamma, x' - z'\pi, z) \quad \text{for } P_{uvz} \in \mathcal{P}_{uvz}, \pi \in \mathbb{R}^{k \times d}, \gamma \in \mathbb{R}^{k \times 1}.$$

The submodel \mathcal{P}_β is the subset of \mathcal{P} with $\det(\pi'\pi) \neq 0$ and $\gamma \in \text{col}(\pi)$. For $P \notin \mathcal{P}_\beta$ such that $\pi(P) = 0$ and $\gamma(P) = 0$, the set of pertinent paths $\mathcal{P}_{P,\beta}$ consists of paths of the form $dQ_t(x, y, z) = dQ_{t,uvz}(y - z'(t\hat{\pi}_t)\beta_t), x' - z'(t\hat{\pi}_t), z)$ for $Q_{t,uvz}$ in $\mathcal{P}_{uvz}, \hat{\pi}_t \rightarrow \hat{\pi}, \beta_t \rightarrow \beta$, and $\det(\hat{\pi}'\hat{\pi}) \neq 0$. This can be seen by considering a path Q_t toward P such that $[\pi(Q_t) - 0]/t \rightarrow \hat{\pi}$ and $[\gamma(Q_t) - 0]/t \rightarrow \hat{\pi}\beta$. If $\det(\hat{\pi}'\hat{\pi}) = 0$, then there are many paths taking values in \mathcal{P}_β that have different limits of β .

Now, we characterize the scores and derive $\dot{P}_{P,\beta}$ and β_P . Being a path, $Q_{t,uvz}$ has its own model score g_{uvz} .⁸ Note that the only essential restrictions on $Q_{t,uvz}$ are $\int uQ_{t,uvz}(du, dv', z) = 0$ and $\int vQ_{t,uvz}(du, dv', z) = 0$ for almost every z . Therefore,

$$0 = \frac{1}{t} \frac{\int u(Q_{t,uvz} - P_{uvz})(du, dv', z)}{\int P_{uvz}(du, dv', z)} \rightarrow \frac{\int ug_{uvz}P_{uvz}(du, dv', z)}{\int P_{uvz}(du, dv', z)} = \mathbb{E}_P[ug_{uvz} | z].$$

⁷Domination and differentiability are not necessary as long as each path is differentiable in quadratic mean (Van der Vaart (1998, Section 7.2)). We assume this for illustration of explicit derivation of scores.

⁸For example, for a parametric submodel $(u, v') \perp\!\!\!\perp z, (u, v')' \sim N(0, \Sigma), z \sim N(\zeta, I)$ with $[\zeta(Q_{t,uvz}) - \zeta(P)]/t \rightarrow \dot{\zeta}, [\Sigma(Q_{t,uvz}) - \Sigma(P)]/t \rightarrow \dot{\Sigma}$, we have $g_{uvz} = \frac{1}{2}[u \ v']\Sigma(P)^{-1}\dot{\Sigma}\Sigma(P)^{-1}\begin{bmatrix} u \\ v \end{bmatrix} - \frac{1}{2}\text{tr}(\Sigma(P)^{-1}\dot{\Sigma}) + (z - \zeta(P))'\dot{\zeta}$.

Similarly, $\mathbb{E}_P[v g_{uvz} | z] = 0$. Thus, the set of model scores $\dot{\mathcal{P}}_{P,uvz}$ for \mathcal{P}_{uvz} consists of zero-mean functions in the $L_2(P_{uvz})$ -orthocomplement of the set of functions of the form $uf(z)$ and $vf(z)$. With this, the model score for Q_t is given by

$$\begin{aligned} \frac{dQ_t - dP}{t dP} &= \frac{dQ_{t,uvz}(y - z't\dot{\pi}_t\beta_t, x' - z't\dot{\pi}_t, z) - dP_{uvz}(y - z't\dot{\pi}_t\beta_t, x' - z't\dot{\pi}_t, z)}{t dP} \\ &\quad + \frac{dP_{uvz}(y - z't\dot{\pi}_t\beta_t, x' - z't\dot{\pi}_t, z) - dP_{uvz}(y, x', z)}{t dP} \\ &\rightarrow g = g_{uvz}(y, x', z) - z'\dot{\pi}\beta \frac{\partial}{\partial u} dP_{uvz} - z'\dot{\pi} \frac{\partial}{\partial v} dP_{uvz}. \end{aligned}$$

Thus, $\dot{\mathcal{P}}_{P,\beta}$ is the set of g of this form with $\det(\dot{\pi}'\dot{\pi}) \neq 0$. By integration by parts,

$$\mathbb{E}_P[ug | z] = \frac{- \int uz'\dot{\pi}\beta \frac{\partial}{\partial u} dP_{uvz}(du, dv', z) - \int uz'\dot{\pi} \frac{\partial}{\partial v} dP_{uvz}(du, dv', z)}{\int dP_{uvz}(du, dv', z)} = z'\dot{\pi}\beta.$$

Similarly, $\mathbb{E}_P[v'g | z] = z'\dot{\pi}$. Therefore, the limit of β_t is represented, for example, by

$$\beta = (\mathbb{E}_P[zz']^{-1}\mathbb{E}_P[zv'g])^{\rightarrow} (\mathbb{E}_P[zz']^{-1}\mathbb{E}_P[zug]) =: \beta_P(g),$$

where $A^{\rightarrow} := (A'A)^{-1}A'$ denotes the left inverse of A . This map is continuous on $\dot{\mathcal{P}}_{P,\beta}$ and homogeneous of degree zero but nonlinear. Thus, β is weakly regular.

EXAMPLE 2—Nonlinear Regression: Consider the nonlinear regression model

$$y = \pi m(x; \beta) + \varepsilon, \quad \mathbb{E}[\varepsilon | x] = 0,$$

where m is a known function that is continuously differentiable and Lipschitz in β . Assume for ease of exposition that all variables and parameters are scalars. The key identifying assumption is that $\mathbb{E}[y - \pi m(x; \beta) | x] = 0$ uniquely at (π, β) .

Let $\mathcal{P}_{x\varepsilon}$ be the set of distributions of (x, ε) such that $\mathbb{E}[\varepsilon | x] = 0$, $dP_{x\varepsilon}$ is continuously differentiable in ε , and $m(x; b)$ is square-integrable for every b . The model \mathcal{P} on (x, y) is induced by $dP(x, y) = dP_{x\varepsilon}(x, y - \pi m(x; \beta))$ for some $P_{x\varepsilon} \in \mathcal{P}_{x\varepsilon}$. The submodel \mathcal{P}_β is such that $\mathbb{E}[y - \pi m(x; \beta) | x] = 0$ holds uniquely at (π, β) (so $\pi \neq 0$). Pick $P \in \mathcal{P}$ with $\pi(P) = 0$ and consider the paths that induce $\pi(Q_t) = t\dot{\pi}_t$, $\beta(Q_t) = \beta_t$ with $\dot{\pi}_t \rightarrow \dot{\pi} \neq 0$, $\beta_t \rightarrow \beta$. As in Example 1, paths $Q_{t,x\varepsilon}$ in $\mathcal{P}_{x\varepsilon}$ satisfy $\mathbb{E}[\varepsilon g_{x\varepsilon} | x] = 0$:

$$\frac{dQ_t - dP}{t dP} = \frac{dQ_{t,x\varepsilon}(x, y - t\dot{\pi}_t m(x; \beta_t)) - dP(x, y)}{t dP} \rightarrow g = g_{x\varepsilon} - \dot{\pi} m(x; \beta) \frac{\partial}{\partial \varepsilon} dP_{x\varepsilon}.$$

Thus, $\dot{\mathcal{P}}_{P,\beta}$ is the set of g of this form. By integration by parts, $\mathbb{E}_P[\varepsilon g | x] = \dot{\pi} m(x; \beta)$. Therefore, $(\dot{\pi}_P g, \beta_P(g))$ can be given as the minimizer of $\mathbb{E}_P[(\mathbb{E}_P[\varepsilon g | x] - cm(x; b))^2]$ with respect to (c, b) . This is homogeneous of degree zero and continuous for β_P ; hence β is weakly regular.

EXAMPLE 3—Nonlinear GMM: Consider a nonlinear moment condition that identifies a parameter $(\pi, \beta) \in \mathbb{E} \times \mathbb{B} \subset \mathbb{R}^k \times \mathbb{R}^d$,

$$\mathbb{E}[M_i(\pi, \beta)] = m(\pi, \beta) = 0$$

for a random process M_i (e.g., $Z_i h(X_i; c, b)$ for some X_i and Z_i), indexed by (c, b) , and $\ell \geq k + d$. Let \mathbb{D} be the space of Lipschitz functions $m : \mathbb{E} \times \mathbb{B} \rightarrow \mathbb{R}^\ell$, equipped with the Sobolev-type norm $\|m\| := \|m\|_\infty + \|dm/d\pi'\|_\infty$. Let \mathcal{P}_M be the set of probability distributions P_M of zero-mean stochastic processes taking values in \mathbb{D} . The model \mathcal{P} can be represented as the set of distributions P of M_i such that, for a Borel $B \subset \mathbb{D}$, $P(M_i \in B) = P_m(M_i - m \in B)$ for $m \in \mathbb{D}$. The submodel \mathcal{P}_β is the subset of \mathcal{P} whose mean function m is in the subset $\mathbb{D}_{P,\beta}$ of \mathbb{D} of functions that have unique zeros. Recall that we are interested in the paths along which m vanishes in β at rate t , that is, $m_t(c, b) = m_{0,t}(c) + t\dot{m}_t(c, b)$ where $m_{0,t}(\pi_t) = 0$, $m_t(\pi_t, \beta_t) = 0$, $\dot{m}_t \rightarrow \dot{m}$, $m_{0,t} \rightarrow m_0$, $\frac{d}{d\pi} m_{0,t} \rightarrow \frac{d}{d\pi} m_0$, and $m_0(\pi) = 0$. So, we can write $Q_t \xrightarrow{\text{DQM}} P$ in \mathcal{P}_β using a path $Q_{t,M} \xrightarrow{\text{DQM}} P_M$ in \mathcal{P}_M as

$$Q_t(M_i \in B) = Q_{t,M}(M_i - m_t \in B) \quad \text{for Borel } B,$$

where $[m_t - m_0]/t \rightarrow \dot{m} \in \mathbb{D}_{P,\beta}$, $[\pi_t - \pi]/t \rightarrow \dot{\pi}$, and $\beta_t \rightarrow \beta$. Note that $[m_t(\cdot, \cdot) - m_0(\cdot)]/t = \int M(\cdot, \cdot)[dQ_t - dP]/t \rightarrow \mathbb{E}_P[M(\cdot, \cdot)g]$, so the moment function is regular. The moment conditions imply

$$\begin{aligned} 0 &= \frac{\mathbb{E}_{Q_t}[M(\pi_t, \beta_t)] - \mathbb{E}_P[M(\pi, \beta_t)]}{t} \\ &= \frac{\mathbb{E}_{Q_t}[M(\pi_t, \beta_t)] - \mathbb{E}_{Q_t}[M(\pi, \beta)]}{t} \\ &\quad + \int M(\pi, \beta) \frac{dQ_t - dP}{t} \rightarrow \frac{dm_0(\pi)}{d\pi'} \dot{\pi} + \mathbb{E}_P[M(\pi, \beta)g]. \end{aligned}$$

So $(\dot{\pi}_P(g), \beta_P(g)) \in \mathbb{R}^k \times \mathbb{B}$ can be cast as the zero of the RHS. If we replace g by kg for a scalar k , then $(k\dot{\pi}, \beta)$ gives the corresponding zero; therefore, $\dot{\pi}_P$ is homogeneous of degree 1 and β_P of degree zero. However, if $\beta_P(g_1) \neq \beta_P(g_2)$, then we have no reason to expect $\dot{\pi}_P(g_1 + g_2)$ to match $\dot{\pi}_P(g_1) + \dot{\pi}_P(g_2)$; therefore, π is “directionally differentiable” but not regular. This dovetails with the fact that the distribution of $\sqrt{n}(\hat{\pi} - \pi)$ is nonstandard (Stock and Wright (2000)). Nonetheless, β_P is weakly regular.

2.2. Fundamental Impossibility

The utility of our theoretical formalism can be readily harvested in the following theorem. It gives a formal proof to the conventional wisdom that a “weakly identified” parameter cannot be estimated consistently or pivotally (see, inter alia, Staiger and Stock (1997), Stock and Wright (2000), Andrews and Cheng (2012))—but not as a characteristic of a specific estimation method—as a direct consequence of the characteristic of the model.⁹ This result can also be viewed as a generalized proof of nonexistence of a consistent test conjectured by Hahn, Ham, and Moon (2011).¹⁰ Distinct but related are the impossibility

⁹Consistent estimation may be possible in linear IV models if the number of weak instruments tends to infinity and some other conditions are met. In this case, the structural parameter is not weakly regular.

¹⁰Their setup can be translated into ours by taking \mathbb{B} to be the product space for two estimators compared in the Hausman test, observing that a regular parameter is trivially weakly regular.

results by Dufour (1997) and Hirano and Porter (2015); their setup is a generalization of the weak linear IV structure, whereas our setup is a generalization of the weak identification phenomenon. Indeed, Dufour (1997) showed nonexistence of bounded confidence sets (which is “stronger” than nonexistence of consistent estimators) while there exist weakly regular parameters that admit bounded confidence sets;¹¹ Hirano and Porter (2015) showed the impossibility of unbiased estimation while there exist weakly regular parameters that admit unbiased estimation (Andrews and Armstrong (2017)).

THEOREM 2—Impossibility of Consistent and Equivariant Estimation: *There is no consistent sequence of estimators of a nontrivial weakly regular parameter; there is no consistent sequence of nontrivial tests of a nontrivial weakly regular parameter; there is no equivariant-in-law sequence of estimators of a nontrivial weakly regular parameter with a separable limit law.*¹²

REMARK: The first two claims are straightforward given the definition. The third claim exploits the fact that the asymptotic distribution of $\hat{\beta}_n$ is “continuous” in local alternatives (a consequence of Le Cam’s third lemma); since $\beta(Q_n)$ is discontinuous at P , $\hat{\beta}_n - \beta(Q_n)$ is necessarily discontinuous at P , failing to be equivariant.

Impossibility of equivariant estimation implies that the asymptotic distribution of any estimator of a weakly regular parameter, when centered at the true value, is nonpivotal and not consistently estimable. However, it does not preclude the possibility that there exist *test statistics* whose distributions are pivotal or consistently estimable. In fact, almost any reasonable inference procedure would be based on statistics whose asymptotic distributions are known or estimable; hence, the problem of estimation and the problem of inference bear quite distinct aspects when it comes to weakly regular parameters. This is in stark contrast to the classical context of regular parameters, in which efficient estimation and “efficient” inference are closely related to each other.¹³ This separation partly explains the speciality of current literature on inference problems regarding weak identification.

2.3. Underlying Regular Parameters

The idea to analyze weak regularity is that, in many cases, there exists another parameter that is regular and whose local parameter controls the limit behavior of the weakly regular parameter. In the literature, such a parameter is known as the “reduced-form parameter” and is utilized in various robust inference procedures under weak identification.¹⁴ Then, the weakly regular parameter acts by itself as a transformation of the local parameter of some “underlying” regular parameter; in other words, it is sufficient to know the value of (the local parameter of) the underlying regular parameter in order to infer the value of the weakly regular parameter around the point of identification failure. We now formalize this idea.

¹¹A bounded weakly regular parameter trivially admits a bounded confidence set.

¹²There is no known example of nonseparable Borel measures and they are usually put aside in the standard theory of weak convergence (Van der Vaart and Wellner (1996, p. 24)), so the assumption of separability is innocuous. We hereafter treat it as general impossibility of equivariant estimation.

¹³Van der Vaart (1998, Chapter 25) stated that “[s]emiparametric testing theory has little more to offer than the comforting conclusion that tests based on efficient estimators are efficient.”

¹⁴On the other hand, the weakly regular parameter is often referred to as the “structural parameter.”

DEFINITION—Underlying Regular Parameter: Let $\beta : \mathcal{P}_\beta \rightarrow \mathbb{B}$ be weakly regular at $P \in \mathcal{P}$ relative to $\mathcal{S}_{P,\beta}$. The parameter $\psi : \mathcal{P} \rightarrow \mathbb{D}$ is an *underlying (regular) parameter* for β at P relative to \mathcal{S}_P if it is regular at P relative to \mathcal{S}_P and there exists a continuous map $\beta_{P,\psi} : \mathbb{D}_{P,\beta} \rightarrow \mathbb{B}$ that is homogeneous of degree zero such that

$$\beta(Q_t) \longrightarrow \beta_{P,\psi}(\dot{\psi}_P g) \quad \text{for every } Q_t \in \mathcal{S}_{P,\beta},$$

where $\mathbb{D}_{P,\beta} := \{\delta \in \mathbb{D} : \delta = \dot{\psi}_P g \text{ for some } g \in \dot{\mathcal{P}}_{P,\beta}\}$.

REMARK: A global sufficient condition for an underlying regular parameter is that there exists a map from regular ψ to weakly regular β that is continuous on $\psi(\mathcal{P}_\beta) \subset \mathbb{D}$.

REMARK: Cox (2017) defined the reduced-form parameter as a function of the structural parameter. We take the opposite route: the weakly regular parameter approaches a function of (the local parameter of) an underlying regular parameter.

This definition requires that knowing the local parameter of the underlying regular parameter is enough to recover the value of the weakly regular parameter, that is, the reduction of information from knowing g to knowing $\dot{\psi}_P g$ does not impair the ability to discern β in the limit.

As it turns out, it is straightforward to show that an underlying parameter always exists. If we regard the root likelihood ratio $Q \mapsto dQ^{1/2}/dP^{1/2}$ as a parameter, it is trivially an underlying regular parameter for any weakly regular parameter. However, whether there exists an underlying regular parameter that admits \sqrt{n} estimation is a different matter. For this, we need to investigate each model separately.

LEMMA 3—Existence of Underlying Regular Parameter: *Let $\beta : \mathcal{P}_\beta \rightarrow \mathbb{B}$ be weakly regular. Then, there exist a Banach space \mathbb{D} and an underlying regular parameter $\psi : \mathcal{P} \rightarrow \mathbb{D}$ for β .*

Below, we see that the natural parameters that appear in examples constitute underlying regular parameters.

EXAMPLE 1—Linear IV, continued: Define $\psi := (\gamma, \text{vec}(\pi))$ to be the $(k + kd) \times 1$ parameter of “reduced-form coefficients.” We verify that ψ is an underlying regular parameter for β . Recall that $\dot{\gamma} = \mathbb{E}_P[zz']^{-1}\mathbb{E}_P[zug]$ and $\dot{\pi} = \mathbb{E}_P[zz']^{-1}\mathbb{E}_P[zv'g]$, that is, the local parameter of ψ is a continuous linear functional of the score; therefore, ψ is regular with $\dot{\psi}_P g = (\dot{\gamma}, \dot{\pi})$. For $g \in \dot{\mathcal{P}}_{P,\beta}$, we have $\dot{\gamma} = \dot{\pi}\beta$ and $\beta_P(g) = \dot{\pi}^\top \dot{\gamma}$. Therefore, ψ is an underlying regular parameter for β with $\beta_{P,\psi}(\dot{\psi}) = \dot{\pi}^\top \dot{\gamma}$ defined on $\mathbb{D}_{P,\beta} = \{(\dot{\gamma}, \text{vec}(\dot{\pi})) \in \mathbb{R}^{k \times 1} \times \mathbb{R}^{k \times d} : \det(\dot{\pi}'\dot{\pi}) \neq 0, \dot{\gamma} \in \text{col}(\dot{\pi})\}$. In fact, this underlying parameter admits the direct representation $\beta(Q_t) = \pi(Q_t)^\top \gamma(Q_t)$.

There are other choices of the underlying regular parameter. Let π_d and γ_d be the first $d \times d$ submatrix and $d \times 1$ subvector of π and γ . Then $\psi_d := (\gamma_d, \text{vec}(\pi_d))$ is also an underlying regular parameter since $\beta_P(g) = \dot{\pi}_d^\top \dot{\gamma}_d$. The submatrix and subvector can, in fact, be for any combinations of coefficients on k instruments as long as $\det(\dot{\pi}'_d \dot{\pi}_d) \neq 0$ and $\dot{\gamma}_d \in \text{col}(\dot{\pi}_d)$. This is to say that in overidentified linear IV models ($k > d$), there are many natural choices of underlying regular parameters.

EXAMPLE 2—Nonlinear Regression, continued: Let $\mathbb{D} = \{cm(\cdot; b) : c \in \mathbb{R}, b \in \mathbb{B} \subset \mathbb{R}\}$ be the space of functions spanned by $cm(\cdot; b)$. From the form of β_P , let us speculate that $\psi : \mathcal{P} \rightarrow \mathbb{D}$, $\psi(Q) := \mathbb{E}_Q[y | x = \cdot]$, is an underlying regular parameter for β . At the point of identification failure P , we have $\psi(P) \equiv 0$. Along the paths we consider,

$$\frac{\psi(Q_t) - \psi(P)}{t} \longrightarrow \dot{\psi}_P g := \mathbb{E}[\varepsilon g | x = \cdot] = \dot{\pi} m(\cdot; \beta),$$

which shows regularity of ψ . Next, $(\dot{\pi}, \beta)$ can be cast as the minimizer of $\mathbb{E}_P[(\dot{\psi}_P g(x) - cm(x; b))^2]$, which is homogeneous of degree zero and continuous for β (ergo for $\beta_{P, \psi}$). Conclude that ψ is an underlying regular parameter for β . It may seem surprising that π is not a part of ψ , but it is encoded as the scaling factor of ψ .

EXAMPLE 3—Nonlinear GMM, continued: The moment function $m : \mathbb{E} \times \mathbb{B} \rightarrow \mathbb{R}$ is an underlying regular parameter for β . Regularity is verified in the previous section. The equation that defines $(\dot{\pi}_P, \beta_P)$ can be written as $\frac{dm_Q(\pi)}{d\pi} \dot{\pi} + \dot{m}(\pi, \beta) = 0$. Thus, by taking $(\dot{\pi}_{P, m}(\dot{m}), \beta_{P, m}(\dot{m}))$ to be the zero of this (defined on the subset $\mathbb{D}_{P, \beta} \subset \mathbb{D}$ of functions with unique zeros), one sees that the moment function is an underlying regular parameter for β .

3. MINIMAL SUFFICIENT UNDERLYING REGULAR PARAMETERS

The underlying regular parameters are not unique, and different parameters may exhibit different properties. This section develops two desirable properties for underlying parameters. The motivation is analogous to classical semiparametric efficiency. In a model that contains both a parameter of interest β and a nuisance parameter, there is a variation of the data that is informative of β and a variation that is not informative of β . The classical theory extracts “pure” variation for β and makes an estimator depend only thereon. If β is regular, such estimators share a unique efficient distribution by the virtue of differentiability. In other words, the following two observations hold simultaneously: (1) desirable estimators depend only on pure variation; (2) desirable estimators share a unique distribution.

In the context of weakly regular β , (2) is no longer attainable due to Theorem 2, while (1) still is. In particular, we can make an estimator $\hat{\psi}$ depend only on pure variation for β and then construct an estimator of β using $\hat{\psi}$. Such ψ plays the role of extracting pure Gaussian variation relevant to β . In the local expansion, this means that the local parameter $\dot{\psi}_P$ shares the same set of nuisance scores as β_P . Note that these concepts need to be developed in local terms since the nuisance scores depend on the point of local expansion P .

3.1. Nuisance Tangent Spaces

This section defines the space of nuisance scores for a weakly regular parameter. Two points deserve attention. First, a weakly regular parameter is not defined on some probability distributions. Therefore, we do not want to deem a score nuisance if it affects identification of the weakly regular parameter. Second, a weakly regular parameter is not linear in the local expansion. This calls for a way to discuss nuisance-ness for nonlinear maps.

In the classical semiparametric theory, local parameters are linearly related to the score, leading to a very nice use of the theory of linear operators (Bickel et al. (1993)). The following definition extends the key notion to nonlinear maps defined on a cone.

DEFINITION: Let \mathcal{X} be a linear space and \mathcal{Y} a set. For a map $f : A \rightarrow \mathcal{Y}$ defined on a cone A in \mathcal{X} , define the *range* R and *kernel* N by $R(f) := \{y \in \mathcal{Y} : y = f(x) \text{ for some } x \in A\}$ and $N(f) := \{\tilde{x} \in \mathcal{X} : x \pm \tilde{x} \in A \text{ and } f(x \pm \tilde{x}) = f(x) \text{ for every } x \in A\}$.

REMARK: If f is linear and $A = \mathcal{X}$, they reduce to the standard definitions of a range and a kernel for linear maps (Van der Vaart (1998, p. 361); Bickel et al. (1993, p. 417)).

Now we define the nuisance tangent space for β . Any score that does not affect identification or the value of β only contains information about the path that is irrelevant to β , hence nuisance. Since the tangent space $\dot{\mathcal{P}}_P$ is linear, we separate the space into the space spanned by nuisance scores and its orthocomplement. That orthocomplement, by construction, only contains scores relevant to β .

DEFINITION—Nuisance Tangent Space: For a weakly regular parameter $\beta : \mathcal{P}_\beta \rightarrow \mathbb{B}$, call its kernel $N(\beta_P) \subset \dot{\mathcal{P}}_P$ the *nuisance tangent space* for β . Denote by Π_β the projection operator onto $N(\beta_P)^\perp$ in $L_2(P)$.

REMARK: For a regular parameter ψ , the kernel of its local parameter $N(\dot{\psi}_P)$ corresponds to the tangent space for its nuisance parameter (Van der Vaart (1998, p. 369)).

The definition indicates that $\tilde{g} \in N(\beta_P)$ means $g + \tilde{g} \in \dot{\mathcal{P}}_{P,\beta}$ and $\beta_P(g + \tilde{g}) = \beta_P(g)$ for every $g \in \dot{\mathcal{P}}_{P,\beta}$; the first condition is the preservation of identification and the second the preservation of the value of β . That is, the perturbation of P in the direction of \tilde{g} does not affect identification or distinction of β . The flip side is that if $\tilde{g} \notin N(\beta_P)$, then there exists $g \in \dot{\mathcal{P}}_{P,\beta}$ such that either $g + \tilde{g} \notin \dot{\mathcal{P}}_{P,\beta}$ or $\beta_P(g + \tilde{g}) \neq \beta_P(g)$, so we want our statistical procedures to be sensitive to these directions.

This lemma shows that the nuisance tangent space is indeed a linear space.

LEMMA 4: (i) $N(\beta_P)$ is a linear space. (ii) If $P \in \mathcal{P} \setminus \mathcal{P}_\beta$, then $N(\beta_P) \subset \dot{\mathcal{P}}_P \setminus \dot{\mathcal{P}}_{P,\beta}$. (iii) If $P \in \mathcal{P} \setminus \mathcal{P}_\beta$, then $g \in \dot{\mathcal{P}}_{P,\beta}$ implies $\Pi_\beta g \neq 0$.

Now we derive the nuisance tangent spaces and “efficient scores” in examples.

EXAMPLE 1—Linear IV, continued: We show that $N(\beta_P) = \dot{\mathcal{P}}_{P,uvz}$. First, any \tilde{g} that satisfies $\mathbb{E}_P[v\tilde{g} | z] = 0$ and $\mathbb{E}_P[u\tilde{g} | z] = 0$ is in $N(\beta_P)$ since $\mathbb{E}_P[v'(g + \tilde{g}) | z] = \mathbb{E}_P[v'g | z]$ and $\mathbb{E}_P[u(g + \tilde{g}) | z] = \mathbb{E}_P[ug | z]$, implying that g and $g + \tilde{g}$ share the same $\hat{\pi}$ and $\dot{\gamma}$. Therefore, $\dot{\mathcal{P}}_{P,uvz} \subset N(\beta_P)$. Next, take $g_1, g_2 \in \dot{\mathcal{P}}_{P,\beta}$ share the same $\hat{\pi}$ but different β . This means that we have $\mathbb{E}_P[v'g_1 | z] = \mathbb{E}_P[v'g_2 | z] = z'\hat{\pi}$ and $\mathbb{E}_P[ug_1 | z] = z'\hat{\pi}\beta_1 \neq \mathbb{E}_P[ug_2 | z] = z'\hat{\pi}\beta_2$. If $\tilde{g} \in N(\beta_P)$, then $\mathbb{E}_P[v'(g_1 + \tilde{g}) | z] = z'c_1$, $\mathbb{E}_P[u(g_1 + \tilde{g}) | z] = z'c_1\beta_1$, $\mathbb{E}_P[v'(g_2 + \tilde{g}) | z] = z'c_2$, and $\mathbb{E}_P[u(g_2 + \tilde{g}) | z] = z'c_2\beta_2$ for some c_1, c_2 since \tilde{g} does not affect β . Deduce that $\mathbb{E}_P[v'\tilde{g}] = z'(c_1 - \hat{\pi}) = z'(c_2 - \hat{\pi})$ and $\mathbb{E}_P[u\tilde{g} | z] = z'(c_1 - \hat{\pi})\beta_1 = z'(c_2 - \hat{\pi})\beta_2$. Since $\beta_1 \neq \beta_2$, we must have $c_1 = c_2 = \hat{\pi}$. In other words, $\mathbb{E}_P[v'\tilde{g} | z] = 0$ and $\mathbb{E}_P[u\tilde{g} | z] = 0$. Therefore, $g \in \dot{\mathcal{P}}_{P,uvz}$. Conclude that $N(\beta_P) = \dot{\mathcal{P}}_{P,uvz}$.

We note that we can write $g \in \dot{\mathcal{P}}_{P,\beta}$ as the sum of elements in $N(\beta_P)^\perp$ and $N(\beta_P)$. As in Van der Vaart (1998, Example 25.28), $\Pi_\beta g = [z' \dot{\pi} \beta \quad z' \dot{\pi}] \mathbb{E}_P \left[\begin{matrix} u^2 & uv' \\ uv & vv' \end{matrix} \mid z \right]^{-1} \begin{bmatrix} u \\ v \end{bmatrix}$.¹⁵

EXAMPLE 2—Nonlinear Regression, continued: We see that $N(\beta_P) = \dot{\mathcal{P}}_{P,x\varepsilon}$ and $\Pi_\beta g = \dot{\pi} m(x; \beta) \mathbb{E}_P[\varepsilon^2 \mid x]^{-1} \varepsilon$ by the same argument as in Example 1. An interesting observation here is that $\frac{\partial}{\partial \beta} m(x; \beta) \mathbb{E}_P[\varepsilon^2 \mid x]^{-1} \varepsilon$ is in the closure of $\dot{\mathcal{P}}_P$ (but not necessarily in $\dot{\mathcal{P}}_{P,\beta}$). This follows since the linearity of $\dot{\mathcal{P}}_P$ implies that $(m(x; \beta + t) \mathbb{E}_P[\varepsilon^2 \mid x]^{-1} \varepsilon - m(x; \beta) \mathbb{E}_P[\varepsilon^2 \mid x]^{-1} \varepsilon) / t$ is in $\dot{\mathcal{P}}_P$ for every $t > 0$.

EXAMPLE 3—Nonlinear GMM, continued: Characterization of exact $N(\beta_P)$ and $\Pi_\beta g$ requires additional details, but it is clear that $\dot{\mathcal{P}}_{P,m} \subset N(\beta_P)$ since a score \tilde{g} that satisfies $\mathbb{E}_P[M(\cdot, \cdot) \tilde{g}] = 0$ does not affect the equation defining $(\dot{\pi}_P, \beta_P)$. Moreover, if there exist scores \tilde{g} such that $\mathbb{E}_P[M(\pi, \cdot) \tilde{g}]$ is a nonzero constant vector, then they change $\dot{\pi}_P$ but do not change β_P , so are in $N(\beta_P)$. This is the case, for example, when the moment function is separable between π and β . If M is a fully nonlinear function, on the other hand, this is not likely to hold.

3.2. Sufficiency and Minimality of Underlying Regular Parameters

The underlying regular parameters are characterized by the span of their nuisance scores. The first property we want in the underlying regular parameter is that it contain all relevant information about β .

DEFINITION—Sufficiency of Underlying Regular Parameter: Let $\beta : \mathcal{P}_\beta \rightarrow \mathbb{B}$ be weakly regular. An underlying regular parameter $\psi : \mathcal{P} \rightarrow \mathbb{D}$ for β is *sufficient* if $N(\dot{\psi}_P) \subset N(\beta_P)$.

REMARK: A global sufficient condition for sufficiency is that if $Q \in \mathcal{P}_\beta$ and $Q' \in \mathcal{P}$ satisfy $\psi(Q) = \psi(Q')$, then $Q' \in \mathcal{P}_\beta$. In this sense, ψ discerns identification of β .

If $\tilde{\psi}$ is sufficient, then any perturbation that does not affect its local parameter does not affect β . The following example shows that an underlying regular parameter need not be sufficient.

EXAMPLE 1—Insufficiency in Linear IV, continued: Let $d = 1$ and $k > 1$ and consider the underlying regular parameter $\psi_1(Q) = (\gamma_1, \pi_1)$ that induces $\dot{\psi}_{1P} g = (\dot{\gamma}_1, \dot{\pi}_1)$. This parameter uses only the first instrument even though there are more available. Therefore, $N(\dot{\psi}_{1P})$ contains elements $\tilde{g} \propto z_2 \frac{\partial}{\partial \pi} \frac{dP_{uvz}}{dP}$ that only change the value of $\dot{\pi}_2$. However, changing the value of $\dot{\pi}_2$ without changing $\dot{\gamma}_1$, $\dot{\pi}_1$, and $\dot{\gamma}_2$ makes β unidentified; therefore, $g + \tilde{g} \notin \dot{\mathcal{P}}_{P,\beta}$ for $g \in \dot{\mathcal{P}}_{P,\beta}$, so $\tilde{g} \notin N(\beta_P)$. Hence, ψ is not sufficient.

Not surprisingly, the set of all reduced-form coefficients is sufficient.

EXAMPLE 1—Sufficiency in Linear IV, continued: The underlying parameter $\psi(Q) = (\gamma, \text{vec}(\pi))$ is sufficient. To see this, let $\dot{\psi}_P g = (\dot{\gamma}, \text{vec}(\dot{\pi}))$ and $\tilde{g} \in N(\dot{\psi}_P)$. This means

¹⁵For unconditional moment restrictions, $\Pi_\beta g = \mathbb{E}_P[g \begin{pmatrix} u \\ v \end{pmatrix} \otimes z'] \mathbb{E}_P \left[\begin{pmatrix} u^2 & uv' \\ uv & vv' \end{pmatrix} \otimes zz' \right]^{-1} \begin{pmatrix} u \\ v \end{pmatrix} \otimes z = \left(\begin{matrix} \mathbb{E}_P[z z'] \dot{\pi} \beta \\ \text{vec}(\mathbb{E}_P[z z'] \dot{\pi}) \end{matrix} \right)' \mathbb{E}_P \left[\begin{pmatrix} u^2 & uv' \\ uv & vv' \end{pmatrix} \otimes zz' \right]^{-1} \begin{pmatrix} u \\ v \end{pmatrix} \otimes z$.

$\dot{\psi}_P(g + \tilde{g}) = (\dot{\gamma}, \text{vec}(\dot{\pi}))$ for every $g \in \dot{\mathcal{P}}_P$. Therefore, if $g \in \dot{\mathcal{P}}_{P,\beta}$, then $g + \tilde{g} \in \dot{\mathcal{P}}_{P,\beta}$ and $\beta_P(g + \tilde{g}) = \beta_P(g)$, that is, $\tilde{g} \in N(\beta_P)$. Conclude that ψ is sufficient.

REMARK: The contrasting conclusion of Example 1 does not contradict the assumption $\overline{\text{Span } \dot{\mathcal{P}}_{P,\beta}} = \overline{\dot{\mathcal{P}}_P}$ since $(\dot{\pi}\beta, \text{vec}(\dot{\pi}))$ for $\beta \in \mathbb{R}^d$ and nondegenerate $\dot{\pi} \in \mathbb{R}^{k \times d}$ spans the entire space for $(\dot{\gamma}, \text{vec}(\dot{\pi})) \in \mathbb{R}^{d+k \times d}$.

The next property we want is that an underlying regular parameter captures only relevant information for the weakly regular parameter.

DEFINITION—Minimality of Underlying Regular Parameter: Let $\beta : \mathcal{P}_\beta \rightarrow \mathbb{B}$ be weakly regular. An underlying regular parameter $\psi : \mathcal{P} \rightarrow \mathbb{D}$ for β is *minimal* if $N(\beta_P) \subset N(\dot{\psi}_P)$.

REMARK: A global sufficient condition for minimality is that there does not exist a non-injective linear map $f : \mathbb{D} \rightarrow \mathbb{E}$ for some Banach space \mathbb{E} such that $f(\psi)$ is also an underlying regular parameter for β .

Minimality of ψ requires the opposite of sufficiency. For a minimal ψ , any perturbation that does not affect β does not affect ψ . In this sense, a minimal underlying parameter is free of nuisance parameters.

EXAMPLE 1—Minimality in Linear IV, continued: From Example 1 in the previous subsection, $N(\beta_P) = \dot{\mathcal{P}}_{P,uvz}$. Since $N(\beta_P)$ is linear (Lemma 4), for every $g \in \dot{\mathcal{P}}_{P,\beta}$ and $\tilde{g} \in N(\beta_P)$, we have $g + \tilde{g} = (g_{uvz} + \tilde{g}) - z' \dot{\pi} \beta \frac{\partial u}{\partial P} \frac{dP_{uvz}}{dP} - z' \dot{\pi} \frac{\partial v}{\partial P} \frac{dP_{uvz}}{dP}$ and $g_{uvz} + \tilde{g} \in N(\beta_P)$. Thus, $\dot{\psi}_P(g + \tilde{g}) = \dot{\psi}_P(g)$ and $\tilde{g} \in N(\dot{\psi}_P)$ for $\psi = (\gamma, \text{vec}(\pi))$. In other words, $N(\beta_P) \subset N(\dot{\psi}_P)$, implying that ψ is minimal. The above argument applies verbatim to $\psi_1 = (\gamma_1, \pi_1)$, so ψ_1 is minimal as well.¹⁶

REMARK: Minimal sufficiency in our definition is of a parameter, while minimal sufficiency in the context of sufficient statistics is of a statistic.

This theorem ensures that a minimal sufficient underlying parameter exists.

THEOREM 5—Existence of Minimal Sufficient Underlying Regular Parameter: *For every weakly regular parameter, there exists a minimal sufficient underlying regular parameter.*

Minimal sufficiency per se is not strong enough to pin down the underlying parameter uniquely. However, minimal sufficient underlying parameters are almost equivalent to each other in terms of the efficient variation they can extract.

THEOREM 6—Characterization of Minimal Sufficient Underlying Regular Parameter: *Let $\beta : \mathcal{P}_\beta \rightarrow \mathbb{B}$ be weakly regular and $\psi : \mathcal{P} \rightarrow \mathbb{D}$ a sufficient underlying regular parameter for β . Then, ψ is minimal if and only if, for any sufficient underlying regular parameter $\phi : \mathcal{P} \rightarrow \mathbb{E}$ for β on a Banach space \mathbb{E} , there exists a linear map $\tau : \mathbb{E} \rightarrow \mathbb{D}$ such that*

$$\tau(\dot{\phi}_P g) = \dot{\psi}_P g \quad \text{for every } g \in \dot{\mathcal{P}}_P.$$

¹⁶Note that $\psi = (\gamma, \text{vec}(\pi))$ is still minimal even in the homoscedastic model. Homoscedasticity helps simplify efficient estimation, but does not help simplify the semiparametric structure itself.

Let us look at examples of minimal sufficient underlying parameters.

EXAMPLE 1—Linear IV, continued: Without any prior knowledge of instrumental irrelevance, $\psi = (\gamma, \text{vec}(\pi))$ is a minimal sufficient underlying regular parameter.

EXAMPLE 2—Nonlinear Regression, continued: We show that the parameter ψ is minimal and sufficient. If $\tilde{g} \in \dot{\mathcal{P}}_P$ implies $\dot{\psi}_P(g + \tilde{g}) = \dot{\psi}_P g$, then by the formula of $\beta_{P,\psi}$, the value of β_P does not change; hence ψ is sufficient. Minimality of ψ is nontrivial; we show that $\dot{\psi}_P(g + \tilde{g}) \neq \dot{\psi}_P g$ implies $\tilde{g} \notin N(\beta_P)$. From the formula of $\beta_{P,\psi}$, if $\beta_P(g_1 + \tilde{g}) = \beta_P(g_1)$ and $\dot{\psi}_P(g_1 + \tilde{g}) \neq \dot{\psi}_P g_1$, then $\dot{\psi}_P(g_1 + \tilde{g})$ can only be different from $\dot{\psi}_P g_1$ in the value of $\dot{\pi}$; let them be $\tilde{\pi}_1 m(\cdot; \beta_1)$ and $\hat{\pi}_1 m(\cdot; \beta_1)$. However, for another $g_2 \in \dot{\mathcal{P}}_{P,\beta}$ with $\beta_P(g_2) = \beta_2 \neq \beta_1$, $\dot{\psi}_P(g_2 + \tilde{g}) = \hat{\pi}_2 m(\cdot; \beta_2) + \tilde{\pi}_1 m(\cdot; \beta_1)$, which yields (if at all) a value of β different from β_2 . Therefore, $\tilde{g} \notin N(\beta_P)$.

EXAMPLE 3—Nonlinear GMM, continued: We show that the moment function m is sufficient and, in some cases, minimal. Recall that $(\dot{\pi}_P, \beta_P)$ is completely characterized by \dot{m}_P through $\frac{dm_0(\pi)}{d\pi} \dot{\pi}_P(g) + (\dot{m}_P g)(\pi, \beta_P(g)) = 0$. Therefore, if g does not alter $\dot{m}_P g$, $\beta_P(g)$ remains unchanged, showing sufficiency. Or equivalently, we can see this by noting $N(\dot{m}_P) = \dot{\mathcal{P}}_{P,m} \subset N(\beta_P)$ from the previous section. If $\dot{\mathcal{P}}_{P,m} = N(\beta_P)$ (intuitively, if the moment function is an involved nonlinear function), the moment function is minimal.

Given a minimal sufficient underlying regular parameter, the problem of estimation or inference for a weakly regular parameter is translated without loss into a problem for the underlying parameter that is regular.

4. WEAK EFFICIENCY FOR WEAKLY REGULAR PARAMETERS

This section defines a notion of efficiency for the estimators of a weakly regular parameter. The difficulty in defining efficiency is that their asymptotic distributions are nonstandard and nonpivotal (Theorem 2). Just as a weakly regular parameter being locally a nonlinear transformation of an underlying regular parameter, an estimator of a weakly regular parameter is often a nonlinear transformation of the estimator of an underlying regular parameter. Then, even when the estimator of the underlying parameter is Gaussian, its nonlinear transformation can in principle be anything.

A key observation is that if the estimator of the underlying parameter contains noise, its transformation also suffers from unnecessary noise. Our idea to define efficiency lies in consideration of such noise; if the estimator of the weakly regular parameter is asymptotically an appropriate transformation of an efficient estimator of a minimal underlying regular parameter, we call it *weakly efficient*. Here, the base estimator must be efficient for otherwise it is contaminated by noise, and the underlying parameter must be minimal for otherwise its transformation might be affected by the estimator of a nuisance component for β . If the underlying parameter is also sufficient, then the weakly efficient estimator may become efficient under strong identification asymptotics.

It is helpful to draw analogy with classical efficiency on regular parameters. Consider two regular parameters, $\nu \in \mathbb{B}$ and $\psi \in \mathbb{D}_\nu \subset \mathbb{D}$, related through a Hadamard differentiable map $\nu_\psi : \mathbb{D}_\nu \rightarrow \mathbb{B}$ by $\nu = \nu_\psi(\psi)$. Since a differentiable map can be approximated locally by a continuous linear map, we may assume that ν_ψ is continuous linear when two parameters admit consistent estimation. To construct an estimator of ν from an estimator of ψ as $\hat{\nu} = T(\hat{\psi})$, there are two aspects to consider: (1) the efficiency of $\hat{\psi}$ and (2) the desirability

of the map T . Note that when $\hat{\psi}$ takes values in \mathbb{D}_ν , there is little motivation to choose T other than ν_ψ . Then, Van der Vaart (1991a) showed that if one has an efficient estimator $\hat{\psi}$ of ψ that takes values in \mathbb{D}_ν , the plug-in estimator $\nu_\psi(\hat{\psi})$ is efficient for ν .

If, on the other hand, $\hat{\psi}$ takes values in a bigger space \mathbb{D} , we need to consider an optimal choice of T . Consider, for example, the strongly-identified linear IV model with unconditional moment restrictions and overidentification, that is, $\mathbb{E}[z(y - x'\beta)] = 0$ with $k > d$. Rewriting $\mathbb{E}[zz']^{-1}\mathbb{E}[z(y - x'\beta)] = \gamma - \pi\beta = 0$, we can let $\nu = \beta$, $\psi = (\gamma, \pi)$, and $\nu_\psi(\psi) = \pi\gamma$. Note that $\mathbb{B} = \mathbb{R}^d$, $\mathbb{D} = \mathbb{R}^{k \times d} \times \mathbb{R}^k$, and \mathbb{D}_ν is a subspace of \mathbb{D} in which γ is in the column space of π and π is of full column rank. The OLS estimator $(\hat{\gamma}, \hat{\pi})$, which is efficient under unconditional moments, takes values outside of \mathbb{D}_ν in that $\hat{\gamma}$ falls outside the column space of $\hat{\pi}$ with probability 1. Therefore, we have to extend T that supports the bigger space \mathbb{D} . This can be considered as a problem of regressing $\hat{\gamma}$ on $\hat{\pi}$ for which the generalized least squares (GLS) estimation is possible since the variance of the error term $\hat{\gamma} - \hat{\pi}\beta$ can be consistently estimated, yielding an optimally weighted GMM estimator for β .

When β is weakly regular, Theorem 5 guarantees that one can find a minimal sufficient underlying regular parameter ψ with which β is related locally through $\beta_{P,\psi}$. The key difference is that $\beta_{P,\psi}$ is not a linear map; it is a continuous homogeneous-of-degree-zero map. Nevertheless, we can construct an estimator of β from an estimator of ψ as $\hat{\beta} = T(\hat{\psi})$, taking into account the two aspects: (1) the efficiency of $\hat{\psi}$ and (2) the desirability of T . The major consequence we must face, however, is that, since the relation of β and ψ is no longer linear, there are multiple choices of T that are admissible from various perspectives. In fact, even when $\hat{\psi}$ takes values in $\mathbb{D}_{P,\beta}$, it may make sense to consider T other than $\beta_{P,\beta}$, which we illustrate by an example.

In this section, we first define what it means for an estimator to be of the form $T(\hat{\psi})$. Then, we show that this estimator admits improvement via Rao–Blackwellization when an efficient estimator for ψ is available. This allows us to define weak efficiency of an estimator among this class of estimators. Finally, we discuss estimators that are efficient under both strong and weak identification asymptotics.

Throughout this section, we assume that there exists a \sqrt{n} efficient estimator of the minimal sufficient underlying regular parameter. Note, however, that not all regular parameters admit \sqrt{n} consistent or efficient estimators in general.

4.1. Regular Estimators

We focus on estimators of β that are transformations of regular estimators of ψ . First, recall the regular estimator for a regular parameter.

DEFINITION—Regular Estimator for Regular Parameter: A sequence of estimators $\hat{\psi}_n$ for a regular parameter $\psi : \mathcal{P} \rightarrow \mathbb{D}$ is called *regular* at $P \in \mathcal{P}$ relative to \mathcal{P}_P if there exists a tight Borel random element L in \mathbb{D} such that

$$\sqrt{n}(\hat{\psi}_n - \psi(Q_n)) \overset{Q_n}{\rightsquigarrow} L \quad \text{for every } Q_n \in \mathcal{P}_P.$$

This sequence is called (*semiparametric*) *efficient* at P relative to \mathcal{P}_P if it attains the distributional lower bound (denote it by L_ψ) of the convolution theorem.

REMARK: The convolution theorem states that $L = L_\psi + L_\eta$, where L_ψ and L_η are independent tight Borel random elements in \mathbb{D} such that $\Pr(L_\psi \in R(\dot{\psi}_P)) = 1$ and $\delta^* L_\psi \sim$

$N(0, \|\dot{\psi}_P^* \delta^*\|_{L_2(P)}^2)$ for every $\delta^* \in \mathbb{D}^*$ (Van der Vaart (1991a, Theorem 2.1)). This is to say, the asymptotic distribution of any regular estimator of a regular parameter is the sum of a Gaussian variable with covariance being the “ L_2 norm” of the efficient influence map and an independent noise. It is efficient when $L_\eta \equiv 0$.

REMARK: If we center $\hat{\psi}_n$ at $\psi(P)$, then $\sqrt{n}(\hat{\psi}_n - \psi(P)) \rightsquigarrow^{Q_n} \dot{\psi}_P g + L$.

We restrict the class of estimators of β to ones that asymptote to functions of estimators of a minimal underlying parameter. Considering a non-minimal underlying parameter per se is not a problem when such ψ can be estimated efficiently, but we require that the estimator of β depends only on its minimal component asymptotically.

DEFINITION—Regular Estimator for Weakly Regular Parameter: A sequence of estimators $\hat{\beta}_n$ for a weakly regular parameter $\beta : \mathcal{P}_\beta \rightarrow \mathbb{B}$ is called *regular* at $P \in \mathcal{P}$ relative to $\mathcal{P}_{P,\beta}$ if there exist a minimal underlying regular parameter $\psi : \mathcal{P} \rightarrow \mathbb{D}$ for β , a regular sequence of estimators $\hat{\psi}_n$ of ψ , and a function $T : \mathbb{D} \rightarrow \mathbb{B}$ that is $(\dot{\psi}_P g + L)$ -almost everywhere continuous for every $g \in \dot{\mathcal{P}}_{P,\beta}$ such that

$$\hat{\beta}_n = T(\sqrt{n}(\hat{\psi}_n - \psi(P))) + o_P(1) \quad \text{for every } Q_n \in \mathcal{P}_{P,\beta}.$$

The asymptotic distribution of a regular estimator follows straightforwardly from the continuous mapping theorem.

PROPOSITION 7: Let $\hat{\beta}_n = T(\sqrt{n}(\hat{\psi}_n - \psi(P))) + o_P(1)$ be a regular sequence of estimators for a weakly regular parameter $\beta : \mathcal{P}_\beta \rightarrow \mathbb{B}$. Then,

$$\hat{\beta}_n \overset{Q_n}{\rightsquigarrow} T(\dot{\psi}_P g + L).$$

EXAMPLE 1—Linear IV, continued: We verify regularity of 2SLS, optimal IV, GMM, limited information maximum likelihood (LIML), continuously updating GMM (CUE), Fuller, and unbiased (Andrews and Armstrong (2017)) estimators.

Observe that the reduced-form coefficients (γ, π) are regular and the 2SLS can be written as a function of their estimators $\hat{\pi}_n = (Z'Z)^{-1}Z'X$ and $\hat{\gamma}_n = (Z'Z)^{-1}Z'Y$:

$$\hat{\beta}_{2SLS} = (\hat{\pi}'_n(Z'Z)\hat{\pi}_n)^{-1}\hat{\pi}'_n(Z'Z)\hat{\gamma}_n = (\sqrt{n}\hat{\pi}'_n\mathbb{E}[zz']\sqrt{n}\hat{\pi}_n)^{-1}\sqrt{n}\hat{\pi}'_n\mathbb{E}[zz']\sqrt{n}\hat{\gamma}_n + o_P(1).$$

The residual is $o_P(1)$ since $(Z'Z)/n$ converges to $\mathbb{E}[zz']$ in probability under every path. Since $\hat{\pi}_n$ is of full column rank with probability 1, $T(\gamma, \pi) = (\pi'\mathbb{E}[zz']\pi)^{-1}\pi'\mathbb{E}[zz']\gamma$ is continuous $(\hat{\gamma}_n, \hat{\pi}_n)$ -almost everywhere.

Under the conditional moment restrictions $\mathbb{E}[\begin{smallmatrix} y-x'\beta \\ x-\pi'z \end{smallmatrix} \mid z] = 0$, the optimal IV is a $(d + dk) \times (1 + d)$ matrix of the form $C(\begin{smallmatrix} \pi'z \\ I_d \otimes z \end{smallmatrix})\mathbb{E}[\begin{smallmatrix} \varepsilon^2 & \varepsilon v' \\ \varepsilon v & vv' \end{smallmatrix} \mid z]^{-1}$ for any $(d + dk) \times (d + dk)$ full-rank matrix C (Newey (1993)). In fact, we can ignore C and π and use the $(k + dk) \times (k + dk)$ matrix $A(z) := (I_{1+d} \otimes z)\mathbb{E}[\begin{smallmatrix} \varepsilon^2 & \varepsilon v' \\ \varepsilon v & vv' \end{smallmatrix} \mid z]^{-1}$. Note that $A(z)$ cannot be consistently estimated because of ε . The optimal IV estimator $(\hat{\beta}_n, \hat{\pi}_n)$ minimizes

$\mathbb{E}_n[\hat{A}(z) \begin{pmatrix} y-x'\hat{\beta}_n \\ x-\hat{\pi}'_n z \end{pmatrix}] \mathbb{E}_n[\hat{A}(z) \begin{pmatrix} y-x'\hat{\beta}_n \\ x-\hat{\pi}'_n z \end{pmatrix}]$, that is,

$$\begin{aligned} \begin{bmatrix} \hat{\beta}_n \\ \text{vec}(\hat{\pi}_n) \end{bmatrix} &= \mathbb{E}_n \left[\hat{A}(z) \begin{pmatrix} x' \\ I_d \otimes z' \end{pmatrix} \right]^{-1} \mathbb{E}_n \left[\hat{A}(z) \begin{pmatrix} y \\ x \end{pmatrix} \right] \\ &= \left(\mathbb{E}_n [\hat{A}(z)(I_{1+d} \otimes z')]^{-1} \mathbb{E}_n \left[\hat{A}(z) \begin{pmatrix} x' \\ I_d \otimes z' \end{pmatrix} \right] \right)^{-1} \\ &\quad \times \mathbb{E}_n [\hat{A}(z)(I_{1+d} \otimes z')]^{-1} \mathbb{E}_n \left[\hat{A}(z) \begin{pmatrix} y \\ x \end{pmatrix} \right]. \end{aligned}$$

Then, we can think of $\mathbb{E}_n[\hat{A}(z)(I_{1+d} \otimes z')]^{-1} \mathbb{E}_n[\hat{A}(z) \begin{pmatrix} x' \\ I_d \otimes z' \end{pmatrix}]$ as a weighted least squares (WLS) estimator of $[\pi' \ I_{dk}]$ and $\mathbb{E}_n[\hat{A}(z)(I_{1+d} \otimes z')]^{-1} \mathbb{E}_n[\hat{A}(z) \begin{pmatrix} y \\ x \end{pmatrix}]$ as a WLS estimator of $[\text{vec}(\pi)^\gamma]$. Thus, if \hat{A} is a continuous function of an estimator of the reduced-form coefficients, the optimal IV estimator is regular.

For GMM, let W be the weighting matrix. The GMM estimator $\hat{\beta}_{\text{GMM}}$ solves $\min_b [Z'(Y - Xb)]' W [Z'(Y - Xb)]$. Write the objective function as

$$\sqrt{n}(\hat{\gamma}_n - \hat{\pi}_n b)' Z' Z W Z' Z \sqrt{n}(\hat{\gamma}_n - \hat{\pi}_n b).$$

The optimal W under unconditional moment restrictions is $\mathbb{E}[(y - x'\beta)^2 z z']^{-1}$, and its feasible version is $\hat{W}_{2\text{SGMM}} = \mathbb{E}_n[(y - x'\hat{\beta}_{2\text{SLS}})^2 z z']^{-1}$. The expectation involved in W (other than the 2SLS estimator) can be consistently estimated. Being a function of the reduced-form OLS and 2SLS, the two-step GMM is regular.

LIML estimates W assuming homoscedasticity, that is, $\hat{W}_{\text{LIML}}(b) = n(Z'Z)^{-1}/\hat{\sigma}^2(b)$ where $\hat{\sigma}^2(b) = \mathbb{E}_n[(y - x'b)^2]$. Since the second and cross moments of y and x can be consistently estimated, LIML is asymptotically only a function of the OLS estimators of the reduced-form coefficients. Similarly, the continuously updating GMM is regular as it uses $\hat{W}_{\text{CUE}}(b) = \mathbb{E}_n[(y - x'b)^2 z z']^{-1}$, which, again, admits consistent estimation.

For Fuller, let $P := Z(Z'Z)^{-1}Z'$. For a constant C , let $\hat{P}_{\text{Fuller}} := P + (C/n)(I - P)$. The Fuller estimator is then given by

$$\begin{aligned} \hat{\beta}_{\text{Fuller}} &= (X' \hat{P}_{\text{Fuller}} X)^{-1} (X' \hat{P}_{\text{Fuller}} Y) \\ &= (C \mathbb{E}[xx'] + \sqrt{n} \hat{\pi}'_n \mathbb{E}[zz'] \sqrt{n} \hat{\pi}_n)^{-1} (C \mathbb{E}[xy] + \sqrt{n} \hat{\pi}'_n \mathbb{E}[zz'] \sqrt{n} \hat{\gamma}_n) + o_P(1). \end{aligned}$$

Thus, Fuller yields a “weighted combination” of OLS ($C = \infty$) and 2SLS ($C = 0$).

Finally, the unbiased estimator is regular. For simplicity, let $d = 1$ and $k = 1$ and assume that we know $\pi > 0$. Also denote the asymptotic variance of $\sqrt{n}(\hat{\gamma}_n, \hat{\pi}_n)$ by $\begin{pmatrix} \sigma_\gamma^2 & \sigma_{\gamma\pi} \\ \sigma_{\gamma\pi} & \sigma_\pi^2 \end{pmatrix}$. Then,

$$\hat{\beta}_{\text{unbiased}} = \frac{1 - \Phi(\sqrt{n} \hat{\pi}_n / \hat{\sigma}_{\pi,n})}{\hat{\sigma}_{\pi,n} \phi(\sqrt{n} \hat{\pi}_n / \hat{\sigma}_{\pi,n})} \left[\hat{\gamma}_n - \frac{\hat{\sigma}_{\gamma\pi,n}}{\hat{\sigma}_{\gamma,n}^2} \hat{\pi}_n \right] + \frac{\hat{\sigma}_{\gamma\pi,n}}{\hat{\sigma}_{\gamma,n}^2},$$

where σ 's are consistently estimated. This is an example in which it makes sense to consider T other than $\beta_{P,\psi}$ even if $\hat{\psi}_n \in \mathbb{D}_{P,\beta}$ almost surely.

EXAMPLE 2—Nonlinear Regression, continued: The estimator $\hat{\beta}_n$ of β that constitutes a minimizer $(\hat{\pi}_n, \hat{\beta}_n)$ of $\mathbb{E}_n[(y - cm(x; b))^2]$ with respect to (c, b) is regular, provided that the estimator $\hat{\psi}_n := \hat{\pi}_n m(\cdot; \hat{\beta}_n)$ of ψ is regular.

EXAMPLE 3—Nonlinear GMM, continued: Let $\hat{m}_n(\cdot, \cdot) := \mathbb{E}_n[M_i(\cdot, \cdot)]$ be a regular estimator of the moment function. The GMM estimator for (π, β) solves

$$\min_{c,b} \sqrt{n} \hat{m}_n(c, b)' W \sqrt{n} \hat{m}_n(c, b)$$

for some $\ell \times \ell$ positive definite matrix W . Oftentimes W is estimated using initial estimates of (π, β) or updated simultaneously with minimization. In the former case, if \hat{W} is continuous in the initial estimates, then the resulting GMM estimator is regular since minimization is continuous in the given norm. In the latter case, if the whole objective function is continuous, the GMM estimator is regular.

4.2. Weakly Efficient Estimators

We show that for any regular estimator of a weakly regular parameter, there exists another regular estimator that is weakly better in terms of convex loss. A strict improvement is possible unless the estimator is already a nonrandom transformation of an efficient estimator of the underlying parameter. For regular $\tilde{\beta} = S(\tilde{\psi})$, a particular improvement is given as the conditional expectation of $\tilde{\beta}$ conditional on an efficient estimator $\hat{\psi}$, that is, $\mathbb{E}[S(\tilde{\psi}) \mid \hat{\psi}]$. To formalize the efficiency gain, however, we use the normalized expression $\tilde{\beta} = T(\sqrt{n}(\tilde{\psi} - \psi(P)))$ for unknown T and $\psi(P)$.

THEOREM 8—Local Asymptotic Rao–Blackwellization: *Let $\beta : \mathcal{P}_\beta \rightarrow \mathbb{B}$ be weakly regular and $\psi : \mathcal{P} \rightarrow \mathbb{D}$ a minimal underlying regular parameter for β . Let $\tilde{\psi}_n$ be a regular sequence of estimators of ψ and $\tilde{\beta}_n = T(\sqrt{n}(\tilde{\psi}_n - \psi(P))) + o_p(1)$ be a regular sequence of estimators of β . Suppose that an efficient regular sequence of estimators $\hat{\psi}_n$ of ψ exists and $\bar{T}(\delta) := \mathbb{E}[T(\delta + L_\eta/\sqrt{n})]$ exists as a Bochner integral. Then $\hat{\beta}_n := \bar{T}(\sqrt{n}(\hat{\psi}_n - \psi(P)))$ is a better regular estimator than $\tilde{\beta}_n$ in the sense that, for every convex continuous loss function $\ell : \mathbb{B} \rightarrow \mathbb{R}$ such that $\ell(\tilde{\beta}_n - \beta(Q_n))$ and $\ell(\hat{\beta}_n - \beta(Q_n))$ are asymptotically equi-integrable under $Q_n \in \mathcal{P}_{\beta, \beta}$,*¹⁷

$$\liminf_{n \rightarrow \infty} \mathbb{E}_{Q_n, *}[\ell(\tilde{\beta}_n - \beta(Q_n))] - \mathbb{E}_{Q_n}^*[\ell(\hat{\beta}_n - \beta(Q_n))] \geq 0.$$

Note that although T and $\psi(P)$ are not known, we can construct a feasible estimator that is asymptotically equivalent to $\hat{\beta}$. Let $\tilde{\beta} = S(\tilde{\psi})$ where $\tilde{\psi}$ is an inefficient estimator of ψ that is asymptotically normal. Then, the Rao–Blackwellized estimator for β is calculated as follows: (1) compute efficient $\hat{\psi}$ and $\text{Var}(\tilde{\psi} \mid \hat{\psi})$; (2) compute $\hat{\beta} = \mathbb{E}_m[S(\hat{\psi} + e_j) \mid \hat{\psi}]$ where e_1, \dots, e_m are drawn i.i.d. from $N(0, \text{Var}(\tilde{\psi} \mid \hat{\psi}))$. This feasible $\hat{\beta}$ has the same property as $\hat{\beta}$ in Theorem 8.

¹⁷ X_n is asymptotically equi-integrable if $\lim_{M \rightarrow \infty} \limsup_{n \rightarrow \infty} \mathbb{E}^*[\mathbb{1}_{\{|X_n| > M\}}] = 0$ (Van der Vaart and Wellner (1996, p. 421)).

REMARK: Theorem 8 is a kind of admissibility requirement for a convex loss, for example, it does not exclude constant estimators. Unlike admissibility, however, it confines attention to the class of regular estimators while providing an improvement method of Rao–Blackwellization. If $\mathbb{B} = \mathbb{R}$, $\hat{\beta}_n$ first-order stochastically dominates $\tilde{\beta}_n$.

REMARK: Efficiency is usually justified for *subconvex* loss functions (Van der Vaart and Wellner (1996, Theorem 3.11.5)). Theorem 8 is in the same spirit but restricts us to *convex* functions.¹⁸ This difference comes from the fact that our best asymptotic distribution is a nonlinear transformation of Gaussian; there is no symmetry of the distribution we can exploit to accommodate subconvexity.

Now we define our efficiency concept and introduce examples.

DEFINITION—Weak Efficiency for Weakly Regular Parameter: A regular sequence of estimators $\hat{\beta}_n$ for a weakly regular parameter β is *weakly (semiparametric) efficient* at $P \in \mathcal{P}$ relative to $\mathcal{P}_{P,\beta}$ if the involved sequence of estimators for the minimal underlying regular parameter $\hat{\psi}_n$ is efficient.

EXAMPLE 1—Linear IV, continued: Suppose that the reduced-form errors are heteroscedastic and the feasible GLS estimator is available. Then, we can improve many estimators by Theorem 8, including even the optimal IV estimator.

Denote by $(\tilde{\gamma}_n, \tilde{\pi}_n)$ and $(\hat{\gamma}_n, \hat{\pi}_n)$ the OLS and GLS estimators of the reduced-form coefficients. By the efficiency of GLS, $[\frac{\sqrt{n}(\tilde{\gamma}_n - \hat{\gamma}_n)}{\sqrt{n}(\tilde{\pi}_n - \hat{\pi}_n)}] \rightsquigarrow [e_\gamma]$. Note that the asymptotic distributions of OLS and GLS are estimable, and GLS and the noise are independent, so we can consistently estimate the distribution of the noise. Then, the Rao–Blackwellized version of 2SLS takes the form $\bar{T}(\sqrt{n}\hat{\gamma}_n, \sqrt{n}\hat{\pi}_n)$ with

$$\bar{T}(\gamma, \pi) := \mathbb{E}[(\pi + e_\pi)' \mathbb{E}[zz'](\pi + e_\pi)]^{-1} (\pi + e_\pi)' \mathbb{E}[zz'](\gamma + e_\gamma).$$

This is weakly efficient since there is no more noise to Rao–Blackwellize. In practice, this expectation can be computed numerically.

To understand why the optimal IV can be improved, note that it exploits the heteroskedasticity of the structural error ε . However, ε cannot be consistently estimated, while the reduced-form errors (u, v) can be. This is where Theorem 8 finds room for improvement. Since WLS is not as efficient as GLS, we can draw the noise and compute many instances of the optimal IV estimator, $[\widehat{\pi}_{Idk}^\pi]^{-1} [\widehat{\text{vec}(\pi)}^\gamma]$, and take the numerical average to compute the Rao–Blackwellized optimal IV estimator.

LIML is known to have no moment and CUE is suspected to have no moment, hence outside the scope of Theorem 8.

EXAMPLE 2—Nonlinear Regression, continued: The form of the local parameter $\dot{\psi}_P g = \mathbb{E}_P[\varepsilon g \mid x = \cdot]$ implies that ε is an influence function, and the *efficient* influence function is given by $\tilde{\psi}_P = \mathbb{E}_P[\varepsilon^2 \mid x = \cdot]^{-1} \varepsilon$. Therefore, if there exists a consistent estimator for $\mathbb{E}[\varepsilon^2 \mid x]$, then minimizing $\mathbb{E}_n[(y - cm(x; b))^2 / \widehat{\mathbb{E}[\varepsilon^2 \mid x]}]$ yields a more efficient

¹⁸Technically, there is no implication between convexity and subconvexity of a function. In this context, subconvexity can be thought of as roughly weaker.

estimator of ψ than minimizing $\mathbb{E}_n[(y - cm(x; b))^2]$.¹⁹ If x is discrete, then $\mathbb{E}_n[\hat{\varepsilon}^2 | x]$ would yield a consistent estimator for $\mathbb{E}[\varepsilon^2 | x]$ if $\hat{\pi}_n$ is consistent toward 0; if y is binary, then the functional form of $\mathbb{E}[\varepsilon^2 | x]$ is fully determined by $\mathbb{E}[y | x]$, hence estimable; if $\mathbb{E}[\varepsilon^2 | x]$ is smooth, we may use a series estimator as in Newey (1994).²⁰ Given that, we can Rao–Blackwellize the original estimator derived from minimizing $\mathbb{E}[(y - cm(x; b))^2]$.

Nonlinear least squares is used to estimate discrete choice models, for example, to avoid derivative calculation. Our method allows us to improve efficiency in such cases.

EXAMPLE 3—Nonlinear GMM, continued: The first part of our theory enables us to find out if there is any nuisance part in the moment function in each specific model (i.e., if the moment function is minimal). Given that, it is often the case that $\mathbb{E}_n[M(\cdot, \cdot)]$ is an efficient estimator of $\mathbb{E}[M(\cdot, \cdot)]$. Then, there is no noise left to Rao–Blackwellize.

Weak efficiency generalizes classical efficiency through a differentiable map to an almost everywhere continuous map. It is therefore straightforward to construct estimators that are “efficient” under both strong and weak identification asymptotics. If ψ is sufficient and T_n asymptotes to a continuous map under weak regularity and to an efficient differentiable map under regularity, the estimator $\hat{\beta} = T_n(\hat{\psi})$ is weakly efficient under weak regularity and efficient under regularity. For example, the Rao–Blackwellized 2SLS exhibits this property. This is desirable since, often in practice, we do not know which asymptotics is a better approximation to the finite-sample situation in hand. Using such an estimator ensures maximal precision regardless of the “correct” asymptotics.

Finally, note that the point of weak efficiency is to exclude nuisance variation from an estimator, and the concept itself does not pin down a unique efficient distribution. For example, constant estimators or any linear transformations of weakly efficient estimators are weakly efficient. In some cases, it is possible to impose additional restrictions to make a weakly efficient estimator unique. In linear IV with $d = k = 1$ where we know the sign of π , the Rao–Blackwellized unbiased estimator is unique. Since many nonlinear functions of Gaussian means admit unique unbiased estimators, it may be possible to uniquely pin down an unbiased weakly efficient estimator in many models. However, unbiased estimators do not always exist (e.g., in linear IV when the sign of π is not known).

5. SIMULATION OF WEAK EFFICIENCY IN LINEAR IV MODELS

To illustrate weak efficiency, we conduct simulation of linear IV models (Example 1) with overidentified conditional moment restrictions and heteroskedasticity. We consider discrete instruments so that we can estimate the heteroscedastic structure without imposing further assumptions. This enables us to compute the optimal IV and the feasible reduced-form GLS estimators. We focus on two estimators, 2SLS and optimal IV, under weak and strong identification asymptotics.

We let $d = 1$ and $k = 3$ so that 2SLS has a second moment. The sample size is chosen to be $n = 1000$. The instrument z_i is uniformly distributed in $\{-1, 1\}^3$, taking eight distinct combinations. The errors (ε_i, v_i) are drawn from a normal distribution with mean 0 and variance depending on z_i as in Table I; this dependence is determined randomly at the beginning of the simulation. The true parameters are given by $\beta = 1$ and $\pi = (1, 1, 1)'/\sqrt{n}$ under weak identification (weakly regular β) and $\beta = 1$ and $\pi = (1, 1, 1)'$ under strong

¹⁹See Van der Vaart (1998, Example 25.66).

²⁰Note that, unlike Example 1, the structural error ε can be consistently estimated since cm can be.

TABLE I
 HETEROSKEDASTICITY OF $(\varepsilon_i, u_i, v_i)$ GIVEN z_i^a

z'_i	$\text{Var}((\varepsilon_i, u_i, v_i) z_i)$	z'_i	$\text{Var}((\varepsilon_i, u_i, v_i) z_i)$
$(-1, -1, -1)$	$\begin{pmatrix} 7.32 & 4.40 & -2.91 \\ 4.40 & 2.65 & -1.75 \\ -2.91 & -1.75 & 1.16 \end{pmatrix}$	$(-1, -1, 1)$	$\begin{pmatrix} 1.91 & 1.77 & -0.14 \\ 1.77 & 2.19 & 0.43 \\ -0.14 & 0.43 & 0.57 \end{pmatrix}$
$(1, -1, -1)$	$\begin{pmatrix} 8.29 & 3.74 & -4.55 \\ 3.74 & 13.41 & 9.66 \\ -4.55 & 9.66 & 14.21 \end{pmatrix}$	$(1, -1, 1)$	$\begin{pmatrix} 3.83 & -2.49 & -6.32 \\ -2.49 & 1.64 & 4.14 \\ -6.32 & 4.14 & 10.46 \end{pmatrix}$
$(-1, 1, -1)$	$\begin{pmatrix} 3.78 & 3.91 & 0.14 \\ 3.91 & 4.66 & 0.74 \\ 0.14 & 0.74 & 0.61 \end{pmatrix}$	$(-1, 1, 1)$	$\begin{pmatrix} 0.55 & 0.32 & -0.23 \\ 0.32 & 0.20 & -0.12 \\ -0.23 & -0.12 & 0.11 \end{pmatrix}$
$(1, 1, -1)$	$\begin{pmatrix} 8.70 & 3.60 & -5.10 \\ 3.60 & 4.54 & 0.94 \\ -5.10 & 0.94 & 6.03 \end{pmatrix}$	$(1, 1, 1)$	$\begin{pmatrix} 1.22 & 0.45 & -0.77 \\ 0.45 & 0.17 & -0.28 \\ -0.77 & -0.28 & 0.49 \end{pmatrix}$

^aSince $u_i = \varepsilon_i + v'_i\beta$, the matrices are of rank 2.

identification (regular β). Simulation runs for 5,000 iterations. The heteroskedasticity-adjusted concentration parameter $\mathbb{E}[\mathbb{E}[vv' | z]^{-1/2} \pi' z z' \pi \mathbb{E}[vv' | z]^{-1/2}]$ is 0.0075 for weak identification and 7.5484 for strong identification.

To compute Rao–Blackwellization, we must derive the feasible GLS estimator for the reduced-form coefficients. A nontrivial aspect of this is that it consists of multiple equations. We handle this by combining them into one big equation:

$$\begin{bmatrix} Y \\ \text{vec}(X) \end{bmatrix} = \begin{bmatrix} Z & 0 \\ 0 & \mathbb{1}_d \otimes Z \end{bmatrix} \begin{bmatrix} \gamma \\ \text{vec}(\pi) \end{bmatrix} + \begin{bmatrix} U \\ \text{vec}(V) \end{bmatrix} =: \begin{bmatrix} Z & 0 \\ 0 & \mathbb{1}_d \otimes Z \end{bmatrix} \psi + \tilde{U}.$$

Consequently, the variance-covariance matrix of \tilde{U} has some nonzero off-diagonal elements. We estimate it with initial OLS coefficients to compute the feasible GLS estimator for (γ, π) . Since GLS is efficient, by orthogonality we have $\text{Var}(\hat{\psi}_{\text{OLS},n} - \hat{\psi}_{\text{GLS},n} | Z) = \text{Var}(\hat{\psi}_{\text{OLS},n} | Z) - \text{Var}(\hat{\psi}_{\text{GLS},n} | Z)$. With this, we compute the conditional expectation of 2SLS conditional on GLS using 100 draws from $\begin{bmatrix} e_\gamma \\ e_\pi \end{bmatrix} \sim N(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \text{Var}(\hat{\psi}_{\text{OLS},n} - \hat{\psi}_{\text{GLS},n} | Z))$. In particular, the RB 2SLS estimator of β is given by²¹


$$\hat{\mathbb{E}}_e \left[((\hat{\pi}_{\text{FGLS},n} + e_\pi)'(Z'Z)(\hat{\pi}_{\text{FGLS},n} + e_\pi))^{-1} (\hat{\pi}_{\text{FGLS},n} + e_\pi)'(Z'Z)(\hat{\gamma}_{\text{FGLS},n} + e_\gamma) \right],$$

where $\hat{\mathbb{E}}_e$ denotes numerical expectation with respect to (e_γ, e_π) .

Rao–Blackwellization of the optimal IV estimator requires a more elaborate procedure, as the optimal IV estimator involves two levels of noises. The first noise comes from the fact that ε , needed to compute $A(z)$, cannot be consistently estimated; if it is estimated with the 2SLS residuals, then it contains noise due to inefficiency of OLS used in 2SLS. The second noise comes from the fact that the optimal IV estimator is a function of the WLS estimator of $\begin{bmatrix} \pi \\ I_{dk} \end{bmatrix}$ and $\begin{bmatrix} \gamma \\ \text{vec}(\pi) \end{bmatrix}$, where the weights are given by the estimated $A(z)$. We use 50 draws to Rao–Blackwellize the first noise and 100 draws for each of the first noise to Rao–Blackwellize the second.

Figure 1(a) is the histograms of 2SLS and RB 2SLS estimators under weak regularity of β . The vertical dotted line indicates the true value, $\beta = 1$. It shows that the distribution

²¹Note that e_γ and e_π are already denormalized by \sqrt{n} .



(a) Histograms of 2SLS and RB 2SLS estimators.

(b) Histograms of optimal IV and RB optimal IV estimators.


FIGURE 1.—Distributions of 2SLS, optimal IV, and their Rao–Blackwellization *under weak regularity of β (weak identification asymptotics)*. Rao–Blackwellization improves the mean squared errors. Simulated with 1,000 observations and 5,000 iterations. Clusters at the boundaries indicate observations outside the range.

of RB 2SLS is more concentrated than 2SLS. Since Rao–Blackwellization does not affect its mean, both estimators share the same bias. Figure 1(b) is the histograms of optimal IV and RB optimal IV estimators for the same run, in which we observe similar results. To connect these histograms to Theorem 8, we consider two loss functions: the mean squared error (MSE) $\ell(x) = x^2$ and the mean absolute error (MAE) $\ell(x) = |x|$, as summarized in Table II. The MSE of 2SLS decreases from 0.84 to 0.20 after Rao–Blackwellization; the MSE of optimal IV from 1.16 to 0.17. The MAE of 2SLS and optimal IV shows similar drop. We see substantial decrease in the losses in both estimators. LAR (Theorem 8) guarantees that the losses of the RB versions do not exceed those of the original ones, at least asymptotically. In this sense, it is preferable to use a weakly efficient estimator whenever available.

Figure 2 is the histograms of the same estimators under regularity of β . From classical results, we know that optimal IV is efficient and 2SLS is not. We see that both RB optimal IV and RB 2SLS coincide with optimal IV under strong identification asymptotics. This suggests that LAR does not alter an already efficient estimator, while it transforms an inefficient estimator into an efficient one. The condition for this to hold can be understood using the analogy introduced at the beginning of Section 4. For an estimator of the form

TABLE II
THE MSE AND MAE OF 2SLS, OPTIMAL IV, AND THEIR RAO–BLACKWELLIZATIONS UNDER WEAK AND STRONG IDENTIFICATION ASYMPTOTICS

	Weak Identification				Strong Identification			
	2SLS		Optimal IV		2SLS		Optimal IV	
	Plain	RB	Plain	RB	Plain	RB	Plain	RB
MSE	0.841	0.196	1.159	0.174	0.001	0.000	0.000	0.000
MAE	0.633	0.428	0.604	0.394	0.030	0.009	0.009	0.009
Observations	1000		1000		1000		1000	
Noise draws	100		50 × 100		100		50 × 100	
Iterations	5000		5000		5000		5000	



(a) Histograms of 2SLS and RB 2SLS estimators.

(b) Histograms of optimal IV and RB optimal IV estimators.

FIGURE 2.—Distributions of 2SLS, optimal IV, and their Rao–Blackwellization *under regularity of β (strong identification asymptotics)*. Asymptotic distribution of RB 2SLS coincides with the optimal IV. Simulated with 1000 observations and 5,000 iterations. Clusters at the boundaries indicate observations outside the range.

$\hat{\beta} = T_n(\hat{\psi})$, if T_n asymptotes to an almost everywhere continuous map under weak regularity of β and to the optimal T under regularity of β , then $\hat{\beta}$ is weakly efficient under weak identification asymptotics and efficient under strong identification asymptotics. This applies to most of known regular estimators.

Computational time of Rao–Blackwellization does not necessarily parallel computational time of the original estimator. There are two components that contribute to computational burden, T_n and $\hat{\psi}_n$, and Rao–Blackwellization only repeats T_n . In our simulation, therefore, Rao–Blackwellization is done very quick. In fact, the most time-consuming part of our simulation is the computation of the original optimal IV estimator, for which derivation of $\hat{\psi}_n$ requires a loop of matrix operations over observations. In our laptop, one iteration of the simulation (computation of 2SLS, optimal IV, their Rao–Blackwellization, and some auxiliary computation) takes less than 0.3 seconds. From a standpoint of strong identification asymptotics, our RB estimators (or the 2SLS with GLS estimators in place of OLS) give a much faster way to compute efficient estimators than to compute optimal IV.

Note that the conditional moment restrictions, $\mathbb{E}[u_i | z_i] = 0$ and $\mathbb{E}[v_i | z_i] = 0$, play a crucial role in this exercise. OLS is inefficient because of them. Relating thereto, another important assumption is the availability of feasible GLS. A notable example in which the form of heteroskedasticity is known a priori is when y_i is binary and one has a conditional moment restriction, $\mathbb{E}[y_i | x_i] = f(x_i)$; the form of heteroskedasticity is uniquely determined by f as $\mathbb{E}[(y_i - f(x_i))^2 | x_i] = f(x_i) - f(x_i)^2$. If f can be estimated, for example when f belongs to some parametric family $\{f_\theta\}$, one may use feasible GLS with no additional loss of generality. In other linear models with an unknown form of heteroskedasticity, feasible GLS with a nonparametric estimator is available under various assumptions (Robinson (1987), Newey (1994)).

6. CONCLUSION

This paper studies weak identification in semiparametric models and investigates efficient estimation. Weak identification is captured by the notion of *weak regularity*, with which the parameter is approximated by a homogeneous-of-degree-zero map of the score.

This nonlinearity implies impossibility of consistent estimation and inference and equivariant estimation. For each weakly regular parameter, there exists an underlying parameter that is regular and fully characterizes the weakly regular parameter locally. An underlying parameter that is minimal and sufficient shares the same nuisance tangent space with the weakly regular parameter, representing the exact amount of information relevant to the weakly regular parameter.

Decomposing estimation of the weakly regular parameter into estimation of the minimal underlying parameter and its transformation, efficiency is discussed in terms of the noise involved in the estimator of the underlying regular parameter. When the estimator of the underlying parameter is inefficient, we can improve the estimator of the weakly regular parameter by taking expectation conditional on the efficient estimator of the underlying parameter. The estimator is called *weakly efficient* if no further improvement is possible. This exploits the property that an efficient estimator of a regular parameter is “asymptotically sufficient,” hence the name *local asymptotic Rao–Blackwellization*. Simulation of the linear IV model demonstrates that the 2SLS and optimal IV estimators can be improved if the feasible GLS estimator of the reduced-form coefficients is available.

APPENDIX A: PROOFS

PROOF OF LEMMA 1: Since $\dot{\mathcal{P}}_P$ is assumed to be linear, if $g \in \dot{\mathcal{P}}_P$, then $ag \in \dot{\mathcal{P}}_P$ for every $a \in \mathbb{R}$. If g is induced by a path $t \mapsto Q_t$ and $a > 0$, then ag can be induced by the path $t \mapsto Q_{at}$, which is the same path up to a scaled index. Therefore, if $Q_t \in \mathcal{P}_P \setminus \mathcal{P}_{P,\beta}$, then $Q_{at} \in \mathcal{P}_P \setminus \mathcal{P}_{P,\beta}$, implying that if $g \in \dot{\mathcal{P}}_P \setminus \dot{\mathcal{P}}_{P,\beta}$, then $ag \in \dot{\mathcal{P}}_P \setminus \dot{\mathcal{P}}_{P,\beta}$. Being defined as a difference between a linear space and a cone, $\dot{\mathcal{P}}_{P,\beta}$ is a cone. Q.E.D.

PROOF OF THEOREM 2: Let $\beta : \mathcal{P}_\beta \rightarrow \mathbb{B}$ be weakly regular and β_P nonconstant.

The first assertion. Suppose that $\hat{\beta}_n : \mathcal{X}^n \rightarrow \mathbb{B}$ is a consistent sequence of estimators, or even weaker, that there exist two paths $Q_{n1}, Q_{n2} \in \mathcal{P}_{P,\beta}$ inducing $g_1, g_2 \in \dot{\mathcal{P}}_{P,\beta}$ such that $\beta_P(g_1) \neq \beta_P(g_2)$ and $\hat{\beta}_n \xrightarrow{Q_{nj}^*} \beta_P(g_j)$ under each $Q_{nj} \in \{Q_{n1}, Q_{n2}\}$. Define $2\varepsilon := \|\beta_P(g_1) - \beta_P(g_2)\|_{\mathbb{B}}$. Denote by Q_{nj}^n the product measure of Q_{nj} on the product sample space \mathcal{X}^n . By the portmanteau theorem (Van der Vaart and Wellner (1996, Theorem 1.3.4)) and the assumption of convergence in outer probability, $\limsup_{n \rightarrow \infty} Q_{n1}^n(\|\hat{\beta}_n - \beta_P(g_1)\|_{\mathbb{B}}^* \geq \varepsilon) \leq 0$ while $\liminf_{n \rightarrow \infty} Q_{n2}^n(\|\hat{\beta}_n - \beta_P(g_1)\|_{\mathbb{B}}^* \geq \varepsilon) \geq \liminf_{n \rightarrow \infty} Q_{n2}^n(\|\hat{\beta}_n - \beta_P(g_1)\|_{\mathbb{B},*} > \varepsilon) \geq 1$. Therefore, Q_{n2}^n is not contiguous to Q_{n1}^n . Being paths, however, Q_{n2}^n must be contiguous to P^n and P^n to Q_{n1}^n (Van der Vaart and Wellner (1996, Lemma 3.10.11 and Theorem 3.10.9)), hence a contradiction.

The second assertion. Let $H_0 : \beta \in \mathbb{B}_0$ and $H_1 : \beta \in \mathbb{B}_1$ be the null and alternative hypotheses such that \mathbb{B}_0 and \mathbb{B}_1 are nonempty. Suppose that $\phi_n : \mathcal{X}^n \rightarrow [0, 1]$ is a consistent sequence of tests of H_0 of level $\alpha < 1$ so that there exist two paths $Q_{n0}, Q_{n1} \in \mathcal{P}_{P,\beta}$ with $\beta_P(g_0) \in \mathbb{B}_0$ and $\beta_P(g_1) \in \mathbb{B}_1$ such that $\phi_n \xrightarrow{Q_{n0}^*} \alpha$ and $\phi_n \xrightarrow{Q_{n1}^*} 1$. Then by the same reasoning, a contradiction follows.

The third assertion. Let $\hat{\beta}_n$ be an equivariant-in-law sequence of estimators of β with a separable limit law, that is, there exists a fixed separable Borel probability measure L on \mathbb{B} such that $\hat{\beta}_n - \beta(Q_n) \overset{Q_n}{\rightsquigarrow} L$ for every $Q_n \in \mathcal{P}_{P,\beta}$. We derive contradiction by constructing two paths along which β takes distinct values but the likelihood ratio of which converges to 1; this means that $\hat{\beta}$ follows the same distribution in both paths by Le Cam’s third lemma; therefore, $\hat{\beta} - \beta$ must follow different distributions. Pick $g_1, g_2 \in \dot{\mathcal{P}}_{P,\beta}$ such

that $\beta_P(g_1) \neq \beta_P(g_2)$ and denote $\beta_1 := \beta_P(g_1)$ and $\beta_2 := \beta_P(g_2)$. Since $\dot{\mathcal{P}}_{P,\beta}$ is a cone (Lemma 1), ag_1 and ag_2 are also in $\dot{\mathcal{P}}_{P,\beta}$ for every $a > 0$, and by homogeneity we have $\beta_P(ag_j) = \beta_j$. For each positive integer k , take $Q_{nk1}, Q_{nk2} \in \mathcal{P}_{P,\beta}$ to be paths that induce scores g_1/k and g_2/k . Let d_{Q_n} denote the metric that metrizes weak topology on \mathbb{B} under Q_n toward separable limits (Van der Vaart and Wellner (1996, p. 73)). For each k , let n_k be such that for every $n \geq n_k$,

$$\int_{\mathcal{X}} \left[\frac{dQ_{nk1}^{1/2} - dP^{1/2}}{1/\sqrt{n}} - \frac{1}{2} \frac{g_1}{k} dP^{1/2} \right]^2 \vee \int_{\mathcal{X}} \left[\frac{dQ_{nk2}^{1/2} - dP^{1/2}}{1/\sqrt{n}} - \frac{1}{2} \frac{g_2}{k} dP^{1/2} \right]^2 < \frac{1}{k},$$

$$d_{Q_{nk1}}(\hat{\beta}_n - \beta(Q_{nk1}), L) \vee d_{Q_{nk2}}(\hat{\beta}_n - \beta(Q_{nk2}), L) < \frac{1}{k}.$$

Then one can take n'_k so that $n'_k \geq n_k$ and $n'_{k+1} > n'_k$ for every k . Construct two paths Q'_{n1} and Q'_{n2} by $Q'_{nj} = Q_{nk_n j}$ where k_n satisfies $n'_k \leq n < n'_{k+1}$. Then $Q'_{nj} \rightarrow^{\text{DQM}} P$ with scores equal to zero and $\hat{\beta}_n - \beta(Q'_{nj})$ converges weakly to L under Q'_{nj} . Now we want to show that dQ'^n_{n2}/dQ'^n_{n1} converges to 1 and invoke Le Cam's third lemma. For this, we adopt the same proof strategy as Van der Vaart (1998, Theorem 7.2). Observe that $\mathbb{E}_{Q'_{n1}}[n(1 - \frac{dQ'^{1/2}_{n2}}{dQ'^{1/2}_{n1}})^2] \leq \int_{\mathcal{X}} [\frac{dQ'^{1/2}_{n1} - dQ'^{1/2}_{n2}}{1/\sqrt{n}}]^2 \rightarrow 0$. By Taylor's theorem, $\log x^2 = -2(1 - x) - (1 - x)^2 + (1 - x)^2 R(1 - x)$ for $R: \mathbb{R} \rightarrow \mathbb{R}$ such that $R(1 - x) \rightarrow 0$ as $x \rightarrow 1$. Then, $\log \frac{dQ'^n_{n2}}{dQ'^n_{n1}}(X_1, \dots, X_n) = \log(\frac{dQ'^n_{n2}}{dQ'^n_{n1}}(X_1) \cdots \frac{dQ'^n_{n2}}{dQ'^n_{n1}}(X_n)) = \sum_{i=1}^n \log \frac{dQ'^n_{n2}}{dQ'^n_{n1}} = -2 \sum_{i=1}^n W_{ni} - \sum_{i=1}^n W_{ni}^2 + \sum_{i=1}^n W_{ni}^2 R(W_{ni})$ where $W_{ni} := 1 - dQ'^{1/2}_{n2}/dQ'^{1/2}_{n1}(X_i)$. We argue that all three terms converge to zero in probability. Under Q'_{n1} ,

$$\left| \mathbb{E} \sum_{i=1}^n W_{ni} \right| = n \left| 1 - \int \frac{dQ'^{1/2}_{n2}}{dQ'^{1/2}_{n1}} dQ'_{n1} \right| \leq \frac{1}{2} \int \left[\frac{dQ'^{1/2}_{n1} - dQ'^{1/2}_{n2}}{1/\sqrt{n}} \right]^2 \rightarrow 0,$$

$$\text{Var} \left(\sum_{i=1}^n W_{ni} \right) \leq \mathbb{E}[nW_{ni}^2] = \mathbb{E} \left[n \left(1 - \frac{dQ'^{1/2}_{n2}}{dQ'^{1/2}_{n1}} \right)^2 \right] \rightarrow 0.$$

These results imply that the expectation and variance of $\sum W_{ni}$ converge to zero; hence it converges to zero in probability. The second result implies that nW_{ni}^2 converges to zero in mean; by the law of large numbers, $\sum W_{ni}^2$ converges to zero in probability. By Markov's inequality, $\Pr(\max_{1 \leq i \leq n} |W_{ni}| > \varepsilon) \leq n \Pr(|W_{ni}| > \varepsilon) \leq n \Pr(nW_{ni}^2 > n\varepsilon^2) \leq \frac{\mathbb{E}[nW_{ni}^2]}{\varepsilon^2} \rightarrow 0$ for every $\varepsilon > 0$. Thus, $\max_{1 \leq i \leq n} |W_{ni}|$ converges to zero in probability, and so does $\max_{1 \leq i \leq n} |R(W_{ni})|$. Therefore, the third term $\sum W_{ni}^2 R(W_{ni})$ converges to zero in probability. We conclude that dQ'^n_{n2}/dQ'^n_{n1} converges to 1 in probability under Q'_{n1} . Since L is separable, by Slutsky's lemma (Van der Vaart and Wellner (1996, Example 1.4.7)), $(\hat{\beta}_n, \frac{dQ'^n_{n2}}{dQ'^n_{n1}}) \overset{Q'_{n1}}{\rightsquigarrow} (\beta_1 + L, 1)$. By Le Cam's third lemma (Van der Vaart and Wellner (1996, Theorem 3.10.7)), $(\beta_2 + L)(B) = \mathbb{E} \mathbb{1}\{\beta_1 + L \in B\} \mathbb{1} = (\beta_1 + L)(B)$ for every Borel $B \subset \mathbb{B}$, which contradicts $\beta_1 \neq \beta_2$. Q.E.D.

PROOF OF LEMMA 3: Denote by \mathbb{D} the Banach space of P -square integrable functions on \mathcal{X} and define $\psi: \mathcal{P} \rightarrow \mathbb{D}$ by $\psi(Q) = dQ^{1/2}/dP^{1/2}$. Note that ψ is regular with derivative $\dot{\psi}_P: \dot{\mathcal{P}}_P \rightarrow \mathbb{D}$, $\dot{\psi}_P g = g$. Thus, we have $\beta_{P,\psi} = \beta_P$. Q.E.D.

PROOF OF LEMMA 4: (i) Trivially, $0 \in N(\beta_P)$. By definition, $\tilde{g}_1, \tilde{g}_2 \in N(\beta_P)$ implies $\tilde{g}_1 + \tilde{g}_2 \in N(\beta_P)$. Take $\tilde{g} \in N(\beta_P)$ and $a > 0$. Since $\dot{\mathcal{P}}_{P,\beta}$ is a cone (Lemma 1) and β_P is homogeneous of degree zero, $\beta_P(g) = \beta_P(g/a) = \beta_P(g/a + \tilde{g}) = \beta_P(g + a\tilde{g})$ for every $g \in \dot{\mathcal{P}}_{P,\beta}$. This means $a\tilde{g} \in N(\beta_P)$. Therefore, $N(\beta_P)$ is linear. (ii) If $P \in \mathcal{P} \setminus \mathcal{P}_\beta$, then $0 \notin \dot{\mathcal{P}}_{P,\beta}$. Since $g \in N(\beta_P) \cap \dot{\mathcal{P}}_{P,\beta}$ implies $\beta_P(g) = \beta_P(g - g) = \beta_P(0)$, $N(\beta_P) \cap \dot{\mathcal{P}}_{P,\beta}$ must be empty. (iii) If $\Pi_\beta g = 0$, then $g \in N(\beta_P)$, which implies $g \notin \dot{\mathcal{P}}_{P,\beta}$ by (ii). Q.E.D.

PROOF OF THEOREM 5: Let $\mathbb{D} = L_2(P)$ and define $\psi : \mathcal{P} \rightarrow \mathbb{D}$ by $\psi(Q) = 2\Pi_\beta dQ^{1/2} / dP^{1/2}$. Then ψ is regular with the derivative $\dot{\psi}_P : \dot{\mathcal{P}}_P \rightarrow \mathbb{D}$, $\dot{\psi}_P g = \Pi_\beta g$. Note that $\beta_P(g) = \beta_P(\Pi_\beta g)$. This implies that ψ is an underlying regular parameter for β and that $N(\dot{\psi}_P) = N(\beta_P)$, which implies minimal sufficiency of ψ . Q.E.D.

PROOF OF THEOREM 6: *Sufficiency.* Assume that, for any sufficient underlying regular parameter $\phi : \mathcal{P} \rightarrow \mathbb{E}$ for β , there exists a map $\tau : \mathbb{E} \rightarrow \mathbb{D}$ such that $\tau(\dot{\phi}_P g) = \dot{\psi}_P g$ for every $g \in \dot{\mathcal{P}}_{P,\beta}$. This means that $N(\dot{\phi}_P) \subset N(\dot{\psi}_P)$. Take ϕ to be minimal; then $N(\beta_P) = N(\dot{\phi}_P) \subset N(\dot{\psi}_P)$. On the other hand, since ψ is assumed to be a sufficient underlying parameter, we have $N(\beta_P) \supset N(\dot{\psi}_P)$.

Necessity. Assume that $\psi : \mathcal{P} \rightarrow \mathbb{D}$ is a minimal sufficient underlying regular parameter for β . Take $\phi : \mathcal{P} \rightarrow \mathbb{E}$ to be another sufficient underlying regular parameter for β . Then $\beta_{P,\psi}(\dot{\psi}_P g) = \beta_{P,\phi}(\dot{\phi}_P g)$ for every $g \in \dot{\mathcal{P}}_{P,\beta}$ and $N(\dot{\psi}_P) = N(\beta_P) \supset N(\dot{\phi}_P)$. The first property implies $\dot{\psi}_P g \in \beta_{P,\psi}^{-1} \beta_{P,\phi}(\dot{\phi}_P g)$ for every $g \in \dot{\mathcal{P}}_{P,\beta}$. The second property implies that if $\dot{\phi}_P g_1 = \dot{\phi}_P g_2$, then $\dot{\psi}_P g_1 = \dot{\psi}_P g_2$. Conclude that there exists a map $\tau : \mathbb{E}_0 \rightarrow \mathbb{D}$ such that $\dot{\psi}_P g = \tau(\dot{\phi}_P g)$ for $g \in \dot{\mathcal{P}}_0$ where $\mathbb{E}_0 := \dot{\phi}_P(\dot{\mathcal{P}}_{P,\beta})$. Since $\dot{\phi}_P$ and $\dot{\psi}_P$ are linear in g , τ must be linear. Finally, one can extend τ on the whole of \mathbb{E} by letting $\tau(e) := \tau(\Pi_{\mathbb{E}_0} e)$. Q.E.D.

PROOF OF PROPOSITION 7: The claim follows by the extended continuous mapping theorem (Van der Vaart and Wellner (1996, Theorem 1.11.1 and Problem 1.11.1)). Q.E.D.

PROOF OF THEOREM 8: Since expectation carries over continuity, $\hat{\beta}_n$ is regular. Write

$$\begin{aligned} & \mathbb{E}_*[\ell(\tilde{\beta}_n - \beta)] - \mathbb{E}^*[\ell(\hat{\beta}_n - \beta)] \\ &= \mathbb{E}_*[\ell(\tilde{\beta}_n - \beta)] - \mathbb{E}[\ell(T(\dot{\psi}_P g + L_\psi + L_\eta) - \beta)] \\ & \quad + \mathbb{E}[\mathbb{E}[\ell(T(\dot{\psi}_P g + L_\psi + L_\eta) - \beta) - \ell(\bar{T}(\dot{\psi}_P g + L_\psi) - \beta) \mid L_\psi]] \\ & \quad + \mathbb{E}[\ell(\bar{T}(\dot{\psi}_P g + L_\psi) - \beta)] - \mathbb{E}^*[\ell(\hat{\beta}_n - \beta)]. \end{aligned}$$

The first and third differences converge to zero by Proposition 7 and Van der Vaart and Wellner (1996, Theorem 1.11.3); the second term is nonnegative since the conditional expectation is nonnegative by a generalized Jensen’s inequality (To and Yip (1975)). Q.E.D.

APPENDIX B: GENERAL WEAK LINEAR IV MODELS

This section discusses Example 1 where π approaches a rank deficient matrix. We are interested in paths Q_n such that

$$\pi(Q_n) = \pi + \frac{\dot{\pi}}{\sqrt{n}} + o\left(\frac{1}{\sqrt{n}}\right), \quad \beta(Q_n) = \beta + \frac{\dot{\beta}}{\sqrt{n}} + o\left(\frac{1}{\sqrt{n}}\right),$$

$$\gamma(Q_n) = \pi(Q_n)\beta(Q_n) = \gamma + \frac{\dot{\gamma}}{\sqrt{n}} + o\left(\frac{1}{\sqrt{n}}\right) = \pi\beta + \frac{\dot{\pi}\beta + \pi\dot{\beta}}{\sqrt{n}} + o\left(\frac{1}{\sqrt{n}}\right),$$

where π is of deficient rank $\ell < d$ and $\pi(Q_n)$ is of full column rank for each n .

We make use of a few innocuous simplifications to the population model. First, redefine z, γ, π to be $\mathbb{E}_P[zz']^{-1/2}z, \mathbb{E}_P[zz']^{1/2}\gamma, \mathbb{E}_P[zz']^{1/2}\pi$, so that we have $\mathbb{E}_P[zz'] = I$. Next, by the singular value decomposition, we can write $\pi = USV'$ for a $k \times k$ orthogonal matrix U , a $d \times d$ orthogonal matrix V , and a $k \times d$ diagonal matrix S whose first ℓ elements are positive and all others zero. Then, by redefining $z, x, v, \gamma, \pi, \beta$ to be $U'z, V'x, V'v, U'\gamma, U'\pi V, V'\beta$, we can make π equal to S .²² To sum up, $\mathbb{E}_P[zz'] = I$, π is diagonal with its first ℓ elements positive, and the last $(\ell - k) \times (\ell - d)$ submatrix of $\dot{\pi}$ is of full column rank. Henceforth, we adopt the notation:

$$\dot{\pi} = \begin{bmatrix} \dot{\pi}_{11} & \dot{\pi}_{12} \\ \dot{\pi}_{21} & \dot{\pi}_{22} \end{bmatrix} = \begin{bmatrix} \dot{\pi}_1 \\ \dot{\pi}_2 \end{bmatrix}, \quad \beta = \begin{bmatrix} \beta_1 \\ \beta_2 \end{bmatrix}, \quad \dot{\beta} = \begin{bmatrix} \dot{\beta}_1 \\ \dot{\beta}_2 \end{bmatrix}, \quad \gamma = \begin{bmatrix} \gamma_1 \\ 0 \end{bmatrix}, \quad \dot{\gamma} = \begin{bmatrix} \dot{\gamma}_1 \\ \dot{\gamma}_2 \end{bmatrix},$$

where π_{11} is an $\ell \times \ell$ matrix, π_1 is an $\ell \times d$ matrix, and $\beta_1, \dot{\beta}_1, \gamma_1, \dot{\gamma}_1$ are $\ell \times 1$ vectors.

We show that (γ, π) is regular, β is weakly regular, and surprisingly, β_1 is not regular unless $\dot{\pi}_{12} \equiv 0$. Since $\beta_1(Q_n) \rightarrow \beta_1(P) = \pi_{11}(P)^{\rightarrow} \gamma_1(P)$, we see that β_1 is continuous, and as such, trivially weakly regular. As before, the score is of the form

$$g = g_{uvz} - z'(\dot{\pi}\beta + \pi\dot{\beta}) \frac{\partial}{\partial P} dP_{uvz} - z'\dot{\pi} \frac{\partial v}{\partial P} dP,$$

and we have $\mathbb{E}_P[ug | z] = z'(\dot{\pi}\beta + \pi\dot{\beta}) = z'\dot{\gamma}$ and $\mathbb{E}_P[v'g | z] = z'\dot{\pi}$. Thus, we find

$$\dot{\gamma}_P g = \mathbb{E}_P[zz']^{-1} \mathbb{E}_P[zug] = \begin{bmatrix} \dot{\pi}_{11}\beta_1 + \dot{\pi}_{12}\beta_2 + \dot{\beta}_1 \\ \dot{\pi}_{21}\beta_1 + \dot{\pi}_{22}\beta_2 \end{bmatrix}, \quad \dot{\pi}_P g = \mathbb{E}_P[zz']^{-1} \mathbb{E}_P[zv'g],$$

showing regularity of (γ, π) . Moreover, we can rearrange the equality of $\dot{\gamma}_{P,2}$ to write

$$\beta_{P,2}(g) = (\dot{\pi}_{P,22}g)^{\rightarrow} (\dot{\gamma}_{P,2}g - (\dot{\pi}_{P,21}g)\beta_1),$$

which is continuous and homogeneous of degree zero in g . Therefore, β_2 is weakly regular and so is the entire vector β . From the equality of $\dot{\gamma}_{P,1}$,

$$\dot{\beta}_{P,1}(g) = \dot{\gamma}_{P,1}g - (\dot{\pi}_{P,11}g)\beta_1 - (\dot{\pi}_{P,12}g)\beta_{P,2}(g).$$

Since $\dot{\gamma}_P$ and $\dot{\pi}_P$ are linear in g and $\beta_{P,2}$ is homogeneous of degree zero in g , we see that $\dot{\beta}_{P,1}$ is homogeneous of degree 1 in g . However, this is not linear in g unless $\dot{\pi}_{P,12}g = 0$ for every g . This observation is akin to Example 3 where π is directionally differentiable but not regular in general. The expression of $\beta_{P,2}$ indicates that (γ_2, π_2) —not (γ, π) —is a minimal sufficient underlying parameter for β .

Now we show that 2SLS is a regular estimator. Since $\mathbb{E}_P[zz'] = I$, we can write

$$\hat{\beta}_{2SLS} = (\hat{\pi}'\hat{\pi})^{-1}(\hat{\pi}'\hat{\gamma}) + o_P(1) = [I \ \sqrt{n}I]([I \ \sqrt{n}I]\hat{\pi}'\hat{\pi}[I \ \sqrt{n}I])^{-1}[I \ \sqrt{n}I]\hat{\pi}'\hat{\gamma} + o_P(1).$$

²²Note that multiplying an orthogonal matrix to z does not affect $\mathbb{E}_P[zz'] = I$.

Observe that

$$\begin{aligned}
 \begin{bmatrix} I \\ \sqrt{n}I \end{bmatrix} \hat{\pi}' \hat{\pi} \begin{bmatrix} I \\ \sqrt{n}I \end{bmatrix} &= \begin{bmatrix} \hat{\pi}'_{11} \hat{\pi}_{11} + \hat{\pi}'_{21} \hat{\pi}_{21} & \sqrt{n} \hat{\pi}'_{11} \hat{\pi}_{12} + \sqrt{n} \hat{\pi}'_{21} \hat{\pi}_{22} \\ \sqrt{n} \hat{\pi}'_{12} \hat{\pi}_{11} + \sqrt{n} \hat{\pi}'_{22} \hat{\pi}_{21} & n \hat{\pi}'_{12} \hat{\pi}_{12} + n \hat{\pi}'_{22} \hat{\pi}_{22} \end{bmatrix} \\
 &= \underbrace{\begin{bmatrix} \hat{\pi}'_{11} \hat{\pi}_{11} & \sqrt{n} \hat{\pi}'_{11} \hat{\pi}_{12} \\ \sqrt{n} \hat{\pi}'_{12} \hat{\pi}_{11} & n \hat{\pi}'_{12} \hat{\pi}_{12} + n \hat{\pi}'_{22} \hat{\pi}_{22} \end{bmatrix}}_{O_P(1)} + \underbrace{\begin{bmatrix} 0 & \sqrt{n} \hat{\pi}'_{21} \hat{\pi}_{22} \\ \sqrt{n} \hat{\pi}'_{22} \hat{\pi}_{21} & 0 \end{bmatrix}}_{O_P(1/\sqrt{n})} \\
 &\quad + o_P\left(\frac{1}{\sqrt{n}}\right), \\
 \begin{bmatrix} I \\ \sqrt{n}I \end{bmatrix} \hat{\pi}' \hat{\gamma} &= \begin{bmatrix} \hat{\pi}'_{11} \hat{\gamma}_1 + \hat{\pi}'_{21} \hat{\gamma}_2 \\ \sqrt{n} \hat{\pi}'_{12} \hat{\gamma}_1 + \sqrt{n} \hat{\pi}'_{22} \hat{\gamma}_2 \end{bmatrix} = \underbrace{\begin{bmatrix} \hat{\pi}'_{11} \hat{\gamma}_1 \\ \sqrt{n} \hat{\pi}'_{12} \hat{\gamma}_1 \end{bmatrix}}_{O_P(1)} + \underbrace{\begin{bmatrix} 0 \\ \sqrt{n} \hat{\pi}'_{22} \hat{\gamma}_2 \end{bmatrix}}_{O_P(1/\sqrt{n})} + o_P\left(\frac{1}{\sqrt{n}}\right).
 \end{aligned}$$

First, let us focus on the $O_P(1)$ terms. Write

$$H := \begin{bmatrix} H_{11} & H_{12} \\ H_{21} & H_{22} \end{bmatrix} := \begin{bmatrix} \hat{\pi}'_{11} \hat{\pi}_{11} & \sqrt{n} \hat{\pi}'_{11} \hat{\pi}_{12} \\ \sqrt{n} \hat{\pi}'_{12} \hat{\pi}_{11} & n \hat{\pi}'_{12} \hat{\pi}_{12} + n \hat{\pi}'_{22} \hat{\pi}_{22} \end{bmatrix}.$$

By the block matrix inversion formula,

$$H^{-1} = \begin{bmatrix} (H_{11} - H_{12}H_{22}^{-1}H_{21})^{-1} & -(H_{11} - H_{12}H_{22}^{-1}H_{21})^{-1}H_{12}H_{22}^{-1} \\ -(H_{22} - H_{21}H_{11}^{-1}H_{12})^{-1}H_{21}H_{11}^{-1} & (H_{22} - H_{21}H_{11}^{-1}H_{12})^{-1} \end{bmatrix}.$$

Thus, the $O_P(1)$ terms make

$$H^{-1} \begin{bmatrix} \hat{\pi}'_{11} \hat{\gamma}_1 \\ \sqrt{n} \hat{\pi}'_{12} \hat{\gamma}_1 \end{bmatrix} = \begin{bmatrix} (H_{11} - H_{12}H_{22}^{-1}H_{21})^{-1} (\hat{\pi}'_{11} \hat{\gamma}_1 - H_{12}H_{22}^{-1} \sqrt{n} \hat{\pi}'_{12} \hat{\gamma}_1) \\ (H_{22} - H_{21}H_{11}^{-1}H_{12})^{-1} (\sqrt{n} \hat{\pi}'_{12} \hat{\gamma}_1 - H_{21}H_{11}^{-1} \hat{\pi}'_{11} \hat{\gamma}_1) \end{bmatrix} = \begin{bmatrix} \hat{\pi}_{11}^{-1} \hat{\gamma}_1 \\ 0 \end{bmatrix}.$$

Therefore, we need the $O_P(1/\sqrt{n})$ terms to derive the asymptotic distribution of $\hat{\beta}_{2SLS,2}$. They are, by the matrix differentiation formula and the Woodbury matrix identity,

$$\begin{aligned}
 &H^{-1} \begin{bmatrix} 0 \\ \sqrt{n} \hat{\pi}'_{22} \hat{\gamma}_2 \end{bmatrix} - H^{-1} \begin{bmatrix} 0 & \sqrt{n} \hat{\pi}'_{21} \hat{\pi}_{22} \\ \sqrt{n} \hat{\pi}'_{22} \hat{\pi}_{21} & 0 \end{bmatrix} H^{-1} \begin{bmatrix} \hat{\pi}'_{11} \hat{\gamma}_1 \\ \sqrt{n} \hat{\pi}'_{12} \hat{\gamma}_1 \end{bmatrix} \\
 &= H^{-1} \begin{bmatrix} 0 \\ \sqrt{n} \hat{\pi}'_{22} (\hat{\gamma}_2 - \hat{\pi}_{21} \hat{\pi}_{11}^{-1} \hat{\gamma}_1) \end{bmatrix} \\
 &= \begin{bmatrix} -\hat{\pi}_{11}^{-1} \sqrt{n} \hat{\pi}_{12} (n \hat{\pi}'_{22} \hat{\pi}_{22})^{-1} \sqrt{n} \hat{\pi}'_{22} (\hat{\gamma}_2 - \hat{\pi}_{21} \hat{\pi}_{11}^{-1} \hat{\gamma}_1) \\ (n \hat{\pi}'_{22} \hat{\pi}_{22})^{-1} \sqrt{n} \hat{\pi}'_{22} (\hat{\gamma}_2 - \hat{\pi}_{21} \hat{\pi}_{11}^{-1} \hat{\gamma}_1) \end{bmatrix}.
 \end{aligned}$$

In short,

$$\hat{\beta}_{2SLS} = \begin{bmatrix} \hat{\pi}_{11}^{-1} \hat{\gamma}_1 \\ (n \hat{\pi}'_{22} \hat{\pi}_{22})^{-1} \sqrt{n} \hat{\pi}'_{22} (\sqrt{n} \hat{\gamma}_2 - \sqrt{n} \hat{\pi}_{21} \hat{\pi}_{11}^{-1} \hat{\gamma}_1) \end{bmatrix} + o_P(1).$$

Thus, the upper half converges in probability to β_1 and the lower half to a function of $\sqrt{n}\hat{\gamma}_2$, $\sqrt{n}\hat{\pi}_{21}$, and $\sqrt{n}\hat{\pi}_{22}$, showing regularity of 2SLS.²³

Under this asymptotics, therefore, we only need to Rao–Blackwellize with respect to $(\hat{\gamma}_2, \hat{\pi}_{21}, \hat{\pi}_{22})$. However, since the coordinate projection of an efficient estimator is efficient and it does not harm to Rao–Blackwellize with respect to a strongly identified parameter, we see that the Rao–Blackwellized 2SLS for the local-to-zero asymptotics derived in Example 1 in the main text is also weakly efficient under this asymptotics.

REFERENCES

- ANDREWS, D. W. K., AND X. CHENG (2012): “Estimation and Inference With Weak, Semi-Strong, and Strong Identification,” *Econometrica*, 80, 2153–2211. [733,734,740]
- ANDREWS, D. W. K., M. J. MOREIRA, AND J. H. STOCK (2006): “Optimal Two-Sided Invariant Similar Tests for Instrumental Variables Regression,” *Econometrica*, 74, 715–752. [735]
- ANDREWS, I., AND T. B. ARMSTRONG (2017): “Unbiased Instrumental Variables Estimation Under Known First-Stage Sign,” *Quantitative Economics*, 8, 479–503. [741,749]
- ANDREWS, I., AND A. MIKUSHEVA (2016): “A Geometric Approach to Nonlinear Econometric Models,” *Econometrica*, 84, 1249–1264. [733]
- BICKEL, P. J., C. A. J. KLAASSEN, Y. RITOV, AND J. A. WELLNER (1993): *Efficient and Adaptive Estimation for Semiparametric Models*. Baltimore and London: Johns Hopkins University Press. [733,736,744]
- COX, G. F. (2017): “Advances in Weak Identification and Robust Inference for Generically Identified Models,” Ph.D. thesis, Yale University. New Haven. [742]
- DUFOUR, J.-M. (1997): “Some Impossibility Theorems in Econometrics With Applications to Structural and Dynamic Models,” *Econometrica*, 65, 1365–1387. [733,735,741]
- HAHN, J., J. C. HAM, AND H. R. MOON (2011): “The Hausman Test and Weak Instruments,” *Journal of Econometrics*, 160, 289–299. [740]
- HIRANO, K., AND J. R. PORTER (2015): “Location Properties of Point Estimators in Linear Instrumental Variables and Related Models,” *Econometric Reviews*, 34, 719–732. [733,741]
- KAJI, T. (2021): “Supplement to ‘Theory of Weak Identification in Semiparametric Models,’” *Econometrica Supplemental Material*, 89, <https://doi.org/10.3982/ECTA16413>. [733]
- MÜLLER, U. K. (2011): “Efficient Tests Under a Weak Convergence Assumption,” *Econometrica*, 79, 395–435. [735]
- MÜLLER, U. K., AND Y. WANG (2019): “Nearly Weighted Risk Minimal Unbiased Estimation,” *Journal of Econometrics*, 209, 18–34. [735]
- NEWKEY, W. K. (1993): “Efficient Estimation of Models With Conditional Moment Restrictions,” in *Handbook of Statistics*, Vol. 11, ed. by G. S. Maddala, C. R. Rao, and H. D. Vinod. Amsterdam: North-Holland, 419–454. [749]
- (1994): “Series Estimation of Regression Functionals,” *Econometric Theory*, 10, 1–28. [753,756]
- ROBINSON, P. M. (1987): “Asymptotically Efficient Estimation in the Presence of Heteroskedasticity of Unknown Form,” *Econometrica*, 55, 875–891. [756]
- STAIGER, D., AND J. H. STOCK (1997): “Instrumental Variables Regression With Weak Instruments,” *Econometrica*, 65, 557–586. [733-735,738,740]
- STOCK, J. H., AND J. H. WRIGHT (2000): “GMM With Weak Identification,” *Econometrica*, 68, 1055–1096. [733,734,740]
- TO, T.-O., AND K. W. YIP (1975): “A Generalized Jensen’s Inequality,” *Pacific Journal of Mathematics*, 58, 255–259. [759]
- VAN DER VAART, A. W. (1991a): “Efficiency and Hadamard Differentiability,” *Scandinavian Journal of Statistics*, 18, 63–75. [748,749]
- (1991b): “On Differentiable Functionals,” *Annals of Statistics*, 19, 178–204. [736]

²³Note that regularity here is as an estimator of a weakly regular parameter. While β_1 is weakly regular, it is not regular in general. Therefore, it does not mean that $\sqrt{n}(\hat{\beta}_1 - \beta_1)$ is equivariant. In fact, its asymptotic distribution is the limit of $\sqrt{n}(\hat{\pi}_{11}^{-1}\hat{\gamma}_1 - \beta_1) - \hat{\pi}_{11}^{-1}\sqrt{n}\hat{\pi}_{12}\hat{\beta}_2$, which is not equivariant even under $\pi_{12} = 0$. If we know $\pi_{12} = 0$, then running 2SLS with regressors x_1 and instruments z_1 , for example, yields an equivariant estimator for β_1 .

- (1998): *Asymptotic Statistics*. Cambridge: Cambridge University Press. [733,736,738,741,744,745,753,758]
- VAN DER VAART, A. W., AND J. A. WELLNER (1996): *Weak Convergence and Empirical Processes: With Applications to Statistics*. New York: Springer. [736,741,751,752,757-759]

Co-editor Ulrich K. Müller handled this manuscript.

Manuscript received 6 June, 2018; final version accepted 10 August, 2020; available online 3 September, 2020.