

Efficient Tests under a Weak Convergence Assumption*

Ulrich K. Müller
Princeton University
Economics Department
Princeton, NJ, 08544

Abstract

The paper studies the asymptotic efficiency and robustness of hypothesis tests when models of interest are defined in terms of a weak convergence property. The null and local alternatives induce different limiting distributions for a random element, and a test is considered robust if it controls asymptotic size for all data generating processes for which the random element has the null limiting distribution. Under weak regularity conditions, asymptotically robust and efficient tests are then simply given by efficient tests of the limiting problem—that is, with the limiting random element assumed observed—evaluated at sample analogues. These tests typically coincide with suitably robustified versions of optimal tests in canonical parametric versions of the model. This paper thus establishes an alternative and broader sense of asymptotic efficiency for many previously derived tests in econometrics, such as tests for unit roots, parameter stability tests and tests about regression coefficients under weak instruments.

JEL classification: C12, C14

Keywords: robustness, unit root test, semiparametric efficiency

*A previous version of this paper was circulated under the title "An Alternative Sense of Asymptotic Efficiency". The author would like to thank Whitney Newey, three anonymous referees, Mark Watson, seminar participants at Berkeley, Brown, Harvard/MIT, Penn State, Stanford and conference participants at the NBER Summer Institute 2008 and the ESEM Summer Meeting 2008 for helpful comments and discussions, as well as Jia Li for excellent research assistance. Financial support by the NSF through grant SES-0751056 is gratefully acknowledged.

1 Introduction

A continued focus of the recent econometrics literature has been the development of asymptotically optimal inference procedures for nonstandard problems: For instance, Elliott, Rothenberg, and Stock (1996) (abbreviated ERS in the following), Elliott (1999), Müller and Elliott (2003) and Müller (2007) derive optimal tests for an autoregressive unit root in a univariate framework; Elliott and Jansson (2003) derive optimal tests for a unit root with stationary covariates; Elliott, Jansson, and Pesavento (2005) derive optimal tests for the null hypothesis of no cointegration with known cointegrating vector; Jansson (2005) derives optimal tests for the null hypothesis of cointegration; Nyblom (1989), Andrews and Ploberger (1994) and Elliott and Müller (2006) derive optimal tests of parameter stability; Stock and Watson (1996) and Jansson and Moreira (2006) derive optimal inference in regression models with nearly integrated regressors; and Andrews, Moreira and Stock (2006, 2008) derive optimal tests for regression coefficients in the presence of weak instruments. By construction, these tests are optimal for a specific parametric version of the model, usually assuming i.i.d. Gaussian disturbances, in the sense of maximizing local asymptotic power. Furthermore, appropriate versions of these tests are robust in the sense that they yield the same asymptotic rejection probability under the null hypothesis (and local alternatives) for a wide range of data generating processes. The assumption of i.i.d. Gaussian disturbances is thus a natural starting point for the development of asymptotically efficient and robust tests.

Nevertheless, with a focus on efficiency, it is natural to ask whether there exist tests that are as good in the Gaussian case, but have higher local asymptotic power for non-Gaussian versions of the models. And indeed, Jansson (2008) draws on and extends the theory of semiparametrically efficient tests to derive such tests for the unit root null hypothesis in the AR(1) model with i.i.d. driving errors of unknown distribution.¹ Also, as for the tests derived under Gaussianity, Jansson (2008) shows that suitably modified versions continue to have correct asymptotic rejection probability under the null hypothesis for a range of serial correlation structures.

This paper also considers the construction of asymptotically efficient tests for nonstandard problems, but with a stronger focus on robustness. For many models and hypothesis

¹Also see Rothenberg and Stock (1997) on tests of the unit root hypothesis under non-Gaussian disturbances.

tests of interest, typical data generating processes imply the weak convergence of some random element to a limiting random element, whose distribution is different under the null and local alternative. For instance, under the null hypothesis of a unit root, the data (suitably scaled) converges to a Wiener process, and it converges to an Ornstein-Uhlenbeck process under the usual local-to-unity alternative. Typically, it is relatively straightforward to think about efficient tests in the 'limiting problem', where the limiting random element is directly observed. For instance, by the Neyman Pearson Lemma, the Radon-Nikodym derivative of the distribution of an Ornstein-Uhlenbeck process (for some fixed mean reversion parameter) and the distribution of a Wiener process is the best point-optimal test statistics in the limiting version of the unit root testing problem. Now suppose one is sufficiently unsure about the nature of the short run dynamics that one would like the test not to overreject whenever the data converges to a Wiener process—or, more generally, whenever the random element converges to its null limiting distribution. If one restricts attention to tests that are robust in this sense, then it is shown that (under mild regularity conditions), the best test statistic is simply given by the best test in the limiting problem, evaluated at sample analogues. In the unit root testing example, this test is asymptotically equivalent to the best unit root test under Gaussian i.i.d. disturbances, so that the test derived by ERS is this best robust test. Any test that has higher asymptotic power than this test for some non-Gaussian version of the model (such as Jansson's (2008) test) necessarily lacks robustness: its asymptotic rejection probability is larger than the nominal level for some model whose suitably scaled data converges weakly to a Wiener process.

The upshot of this analysis is straightforward: to determine the asymptotically efficient robust test in the sense described above, one only needs to consider efficient tests in the limiting problem, where the limiting random element is assumed observed. The potentially complicated small sample testing problem is thus replaced by a (typically) much simpler one. This aspect of the approach makes it akin to LeCam's Limits of Experiments—see Shiryaev and Spokoiny (2000) and chapter 9 of van der Vaart (1998) for introductions. The arguments, however, have distinct starting points: The Limit of Experiments approach considers a sequence of fully specified models, and derives implications from the limiting behavior of the likelihood ratio process; the approach here, in contrast, defines models in terms of a weak convergence property, and studies efficiency by considering the implied asymptotic properties of tests. Also, what is called the limiting problem here is not the

limit experiment in the sense of LeCam: there may exist inference procedures that are asymptotically better in some versions of the model than the best procedure in the limiting problem (again, think of Jansson’s (2008) unit root test), whereas one key result of the limits of experiments approach is that best procedures in the limit experiment provide an upper bound on what can be achieved in original problem in the limit. Rather, efficient tests in our limiting problem provide an upper bound on the asymptotic performance of *robust* tests in the original problem (with robustness defined as described above).

When the parameter of interest is finitely dimensional, the result of this paper can be described as a particular semiparametric efficiency claim, where the definition of the models of interest in terms of a weak convergence property constitute its nonparametric aspect. In contrast, in the traditional semiparametric efficiency literature, the model is described through moment and other conditions, and the sample consists of many (often i.i.d.) observations from the same model. To date, most analyses of traditional semiparametric efficiency concern the class of models with a locally asymptotically normal (LAN) log-likelihood process. In Section 2.5 below, we explicitly compare the results and underlying logic between our approach and the traditional semiparametric efficient tests in a LAN model, and find them to be quite similar. Beyond the potential appeal of starting with a weak convergence assumption for some problems, the advantage of the approach here is that it yields a general and easily applicable efficiency result also in nonstandard testing problems.

The basic result holds quite generally, including for cases where tests are restricted to satisfy some asymptotic unbiasedness or conditional similarity constraint, or to be invariant. Standard arguments from testing theory can then be applied to identify an efficient test in the limiting problem, possibly under an unbiasedness, conditional similarity or invariance constraint. The results of this paper may thus be seen as a formalization and generalization of Sowell (1996), who employs the Neyman-Pearson Lemma to the limiting measure of partial sums of sample moment conditions to derive GMM parameter stability tests. Additional potential applications include a broader sense of asymptotic efficiency of the test statistics derived in the 14 papers cited in the first paragraph of this introduction. Finally, the results of this paper also imply efficiency of some recent nonstandard methods that take a weak convergence assumption as their starting point, such as those suggested in Müller and Watson (2008a, 2008b) and Ibragimov and Müller (2007).

The remainder of the paper is organized as follows. Section 2 introduces the formal

framework and contains the main result. Section 3 discusses extensions regarding consistently estimable nuisance parameters, invariance restrictions and uniformity issues. A running example throughout Sections 2 and 3 is the problem of testing for an autoregressive unit root in a univariate time series. Section 4 discusses in detail three applications: Elliott and Jansson's (2003) point-optimal tests for unit roots with stationary covariates; Andrews, Moreira and Stock's (2006) optimal tests for linear instrumental variable regressions; and Sowell's (1996) tests for GMM parameter stability. Section 4 concludes. All proofs are collected in an appendix.

2 Efficiency and Robustness under a Weak Convergence Assumption

2.1 Set-up

The following notation and conventions are used throughout the paper: All limits are taken as $T \rightarrow \infty$. If S_1 is a metric space with metric d_{S_1} , then $\mathfrak{B}(S_1)$ is its Borel σ -algebra. If μ is a probability measure on $\mathfrak{B}(S_1)$, then its image measure under the $\mathfrak{B}(S_1) \setminus \mathfrak{B}(S_2)$ measurable mapping $f : S_1 \mapsto S_2$, where S_1 and S_2 are spaces with metrics d_{S_1} and d_{S_2} , is denoted $f\mu$. If no ambiguity arises, we suppress the dummy variable of integration, that is we write $\int f d\mu$ for $\int f(x) d\mu(x)$. By default, the product space $S_1 \times S_2$ is equipped with the metric $d_{S_1} + d_{S_2}$. We write $\mu_T \rightsquigarrow \mu_0$ or $X_T \rightsquigarrow X_0$ for the weak convergence of the random elements X_0, X_1, \dots with probability measures μ_0, μ_1, \dots on $\mathfrak{B}(S_1)$, and $X_T \xrightarrow{p} X$ for convergence in probability. The $\mathbb{R} \mapsto \mathbb{R}$ function $x \mapsto \lfloor x \rfloor$ is the integer part of x .

In a sample of size T , suppose we observe data $Y_T \in \mathbb{R}^{nT}$, which is the T th row of a double-array of $n \times 1$ random vectors. The distribution of Y_T depends on the statistical model m with parameter $\theta \in \Theta$, where Θ is a metric space.² The hypotheses of interest are

$$H_0 : \theta \in \Theta_0 \quad \text{against} \quad H_1 : \theta \in \Theta_1 \quad (1)$$

where $\Theta_0 \cup \Theta_1 = \Theta$.

Let h_T be a sequence of measurable functions $h_T : \mathbb{R}^{nT} \mapsto S$, where S is a complete and separable metric space. Denote by $P_T(m, \theta)$ the distribution of $X_T = h_T(Y_T)$ in model m

²By definition of a statistical model, for any m , the distribution $F_T(m, \theta)$ of Y_T is a probability kernel, i.e. $\int_A F_T(m, \theta)$ is a $\mathfrak{B}(\Theta)$ -measurable function of θ for each $A \in \mathfrak{B}(\mathbb{R}^{nT})$.

with parameter θ , that is $P_T(m, \theta) = h_T F_T(m, \theta)$. Suppose the typical model m satisfies the following weak convergences under the null and alternative hypothesis

$$P_T(m, \theta) \rightsquigarrow P(\theta) \quad \text{pointwise for all } \theta \in \Theta_0 \quad (2)$$

$$P_T(m, \theta) \rightsquigarrow P(\theta) \quad \text{pointwise for all } \theta \in \Theta_1 \quad (3)$$

to some statistical model $P(\theta)$. We assume that the probability measures $P(\theta_1)$ and $P(\theta_2)$ on $\mathfrak{B}(S)$ are absolutely continuous for all $\theta_1, \theta_2 \in \Theta$, so that the measures $P(\theta)$ are dominated by $P(\theta_1)$ for any fixed $\theta_1 \in \Theta$. The parameter θ should be thought of as describing *local* alternatives, such as the magnitude of the Pitman drift.

Unit Root Test Example: Consider testing for a unit root in a model with no deterministic components against the local-to-unity alternative: We observe data $Y_T = (u_{T,1}, \dots, u_{T,T})'$ from the model $u_{T,t} = \rho_T u_{T,t-1} + \nu_{T,t}$ and $u_{T,0} = 0$ for all T , where $\rho_T = 1 - c/T$ for some fixed $c \geq 0$, and the hypotheses are $H_0 : c = 0$ against $H_1 : c > 0$ (so that $n = 1$, $\theta = c$, $\Theta_0 = \{0\}$ and $\Theta_1 = (0, \infty)$). Let $\hat{\omega}_T^2$ be a specific, "reasonable" long-run variance estimator. With $S = D_{[0,1]}$ the space of cadlag functions on the unit interval, equipped with the Billingsley (1968) metric, a typical model m for the disturbances $\nu_{T,t}$ satisfies $T^{-1/2} \hat{\omega}_T^{-1} u_{T, \lfloor \cdot T \rfloor} = h_T(Y_T) = \hat{J}_T(\cdot) \rightsquigarrow J_c(\cdot)$ on $D_{[0,1]}$, where J_c is an Ornstein-Uhlenbeck process $J_c(s) = \int_0^s e^{-c(s-r)} dW(r)$ with W a standard Wiener process. It is well known that the measure of J_c is absolutely continuous with respect to the measure of $J_0 = W$. \blacktriangle

2.2 Limiting Problem

In the typical model m , the observed random element X_T satisfies $X_T \rightsquigarrow X \sim P(\theta)$. It will be useful in the sequel to first consider in detail the "limiting problem", where X is assumed observed:

$$H_0^{lp} : X \sim P(\theta), \quad \theta \in \Theta_0 \quad \text{against} \quad H_1^{lp} : X \sim P(\theta), \quad \theta \in \Theta_1. \quad (4)$$

Possibly randomized tests of H_0^{lp} are measurable functions $\varphi_S : S \mapsto [0, 1]$, where $\varphi_S(x)$ indicates the probability of rejection conditional on observing $X = x$, so that the overall rejection probability of the test φ_S when $X \sim P(\theta)$ is given by $\int \varphi_S dP(\theta)$. The test φ_S is thus of level α when $\sup_{\theta \in \Theta_0} \int \varphi_S dP(\theta) \leq \alpha$. In many nonstandard problems, no uniformly most powerful test exists, so consider tests that maximize a weighted average power criterion

$$\text{WAP}(\varphi_S) = \int \left(\int \varphi_S dP(\theta) \right) dw(\theta), \quad (5)$$

where w is a probability measure on Θ_1 . In general, the weighting function w describes the importance a researcher attaches to the ability of the test to reject for certain alternatives. A point-optimal test is a special case of a weighted average power maximizing test for a degenerate weighting function w that puts all mass at one point. Also, if a uniformly most powerful test exists, then it maximizes WAP for all choices for w . The WAP criterion is statistically convenient, since by standard arguments, the WAP maximizing test equivalently maximizes power against the single alternative $H_{1,w}^{lp} : X \sim \int P(\theta)dw(\theta)$.

With the WAP criterion as efficiency measure, efficient level- α tests φ_S^* in the limiting problem (4) thus maximize WAP subject to $\sup_{\theta \in \Theta_0} \int \varphi_S dP(\theta) \leq \alpha$.

Unit Root Test Example, ctd: The weak convergence $\hat{J}_T(\cdot) \rightsquigarrow J_c(\cdot)$ leads to the limiting problem where we observe the continuous time process X , and $H_0^{lp} : X \sim J_0(\cdot)$ against $H_1^{lp} : X \sim J_c(\cdot)$ with $c > 0$. As a weighting function in the WAP criterion, consider a degenerate distribution with all mass at c_1 , so that we consider a point-optimal test, just as ERS. By Girsanov's theorem, the Radon-Nikodym derivative of the distribution of J_{c_1} with respect to the distribution of J_0 , evaluated at X , is given by $L(X) = \exp[-\frac{1}{2}c_1(X(1)^2 - 1) - \frac{1}{2}c_1^2 \int_0^1 X(s)^2 ds]$. Thus, by the Neyman Pearson Lemma, the point-optimal test in the limiting problem is $\varphi_S^*(X) = \mathbf{1}[L(X) > cv]$, where cv solves $P(L(J_0(\cdot)) > cv) = \alpha$. \blacktriangle

When Θ_0 is not a singleton, that is if H_0^{lp} is composite, the derivation of a WAP maximizing test is often more involved. The weighted average power maximizing test under a composite null hypothesis is typically given by the Neyman-Pearson test of $H_{0,\Lambda}^{lp} : X \sim \int P(\theta)d\Lambda(\theta)$ against $H_{1,w}^{lp} : X \sim \int P(\theta)dw(\theta)$, where Λ is the least favorable distribution for θ —see chapter 3.8 of Lehmann and Romano (2005) for discussion. For many problems, however, it is difficult to identify the least favorable distribution Λ . To make further progress, researchers often restrict the class of tests under consideration by additional constraints, and derive the best test in the restricted class. Sometimes, the WAP maximizing test in the restricted class turns out to be uniformly most powerful (that is, maximizes WAP for all weighting functions), so that the issue of how to choose an appropriate weighting function is also avoided by imposing additional constraints.

We discuss invariance as a restriction in Section 3.2 below, and focus here on two other constraints on the tests φ_S . First, consider

$$\int \varphi_S dP(\theta) \geq \pi_0(\theta) \quad \text{for all } \theta \in \Theta_1 \tag{6}$$

for some function $\pi_0 : \Theta_1 \mapsto \mathbb{R}$. The formulation (6) allows for a range of cases: with $\pi_0 = 0$,

(6) never binds; with $\pi_0 = \alpha$, (6) imposes unbiasedness; with π_0 equal to the power of the locally best test for some arbitrarily small neighborhood of Θ_0 , (6) effectively selects φ_S to be the locally best test.

Second, consider a (conditional) similarity constraint of the form

$$\int (\varphi_S - \alpha) f_S dP(\theta) = 0 \quad \text{for all } \theta \in \bar{\Theta}_0 \text{ and } f_S \in \mathcal{F}_S \quad (7)$$

for some $\bar{\Theta}_0 \subset \Theta_0$, which would typically be the intersection of Θ_0 with the closure of Θ_1 , and \mathcal{F}_S some set of measurable and bounded functions $f_S : S \mapsto \mathbb{R}$. With \mathcal{F}_S only containing the zero function, (7) never binds. With \mathcal{F}_S only containing the function that is equal to one, (7) imposes similarity. Finally, suppose $\vartheta : S \mapsto U$ is a measurable function, and \mathcal{F}_S contains all $S \mapsto \mathbb{R}$ functions of the form $f_U \circ \vartheta$, where $f_U : U \mapsto \mathbb{R}$ is continuous and bounded. Then (7) amounts to the restriction that the rejection probability of φ_S for θ on the boundary between the null and alternative hypothesis, conditional on $\vartheta(X)$, is equal to α , so that φ_S is a conditionally similar test.

To sum up, we will refer to level- α tests φ_S^* in the limiting problem (4) as efficient when φ_S^* maximizes weighted average power (5), subject to (6) and (7).

2.3 Robustness

In the original hypothesis testing problem (1) with Y_T observed, tests are measurable functions $\varphi_T : \mathbb{R}^{n_T} \mapsto [0, 1]$, where $\varphi_T(y_T)$ indicates the probability of rejection conditional on observing $Y_T = y_T$. The overall rejection probability of the test φ_T in model m is thus given by $\int \varphi_T dF_T(m, \theta)$.

As in the discussion of the limiting problem, we consider weighted average power as the criterion to measure the efficiency of tests φ_T . In analogy to $\text{WAP}(\varphi_S)$ in the limiting problem (5), define

$$\text{WAP}_T(\varphi_T, m) = \int \int \varphi_T dF_T(m, \theta) dw(\theta).$$

Also, define the asymptotic null rejection probability of the test φ_T in model m as

$$\text{ARP}_0(\varphi_T, m) = \sup_{\theta \in \Theta_0} \limsup_{T \rightarrow \infty} \int \varphi_T dF_T(m, \theta).$$

With these definitions, an asymptotically powerful level- α test φ_T has large $\lim_{T \rightarrow \infty} \text{WAP}_T(\varphi_T, m)$, while $\text{ARP}_0(\varphi_T, m) \leq \alpha$. A reasonable definition of an asymptotically robust test is to impose that $\text{ARP}_0(\varphi_T, m) \leq \alpha$ for a large class of models m .

Now the key premise of this paper is that an interesting set of such models is the set \mathcal{M}_0 of models satisfying (2), i.e. \mathcal{M}_0 collects all data generating processes for Y_T such that $P_T(m, \theta) = h_T F_T(m, \theta) \rightsquigarrow P(\theta)$ for all $\theta \in \Theta_0$. We thus call a test φ_T *robust* if its asymptotic null rejection probability is no larger than the nominal level for all models $m \in \mathcal{M}_0$, that is

$$\sup_{\theta \in \Theta_0} \limsup_{T \rightarrow \infty} \int \varphi_T dF_T(m, \theta) \leq \alpha \quad \text{for all } m \in \mathcal{M}_0. \quad (8)$$

The robustness condition (8) is a point-wise restriction: it is only imposed that for any given $\theta \in \Theta_0$ and $m \in \mathcal{M}_0$, the test φ_T has asymptotic rejection probability of at most α . A stronger, uniform restriction analogous to (8) is discussed in Section 3.3 below.

Ideally, an efficient robust test maximizes $\text{WAP}_T(\varphi_T, m)$ for many alternative models m . In analogy to \mathcal{M}_0 , define \mathcal{M}_1 as the set of alternative models m that satisfy (3).

One might think about the relative gain in robustness of tests satisfying (8) rather than the more standard "correct asymptotic null rejection probability for a wide range of primitive assumptions about disturbances that all imply (2)" in two ways.

On the one hand, one might genuinely worry that the true data generating process happens to be in the set of models that satisfy (2), but the disturbances do not satisfy the primitive conditions. This line of argument then faces the question whether such non-standard data generating processes are plausible. Especially in a time series context, primitive conditions are often quite opaque (could it be that interest rates are not mixing?), so it is often not clear how and with what arguments one would discuss such a possibility. It is probably fair to say, however, that very general forms of sufficient primitive conditions for Central Limit Theorems and alike were derived precisely because researchers felt uncomfortable assuming more restricted (but still quite general) conditions, so one might say that imposing (8) constitutes only one more natural step in this progression of generality.

Unit Root Test Example, ctd: The literature has developed a large number of sufficient conditions on the disturbances $\nu_{T,t}$ that imply $\hat{J}_T(\cdot) \rightsquigarrow J_c(\cdot)$ —see, for instance, McLeish (1974) for a martingale difference sequence framework, Wooldridge and White (1988) for mixing conditions, Phillips and Solo (1992) for linear process assumptions and Davidson (2002) for near-epoch dependence. Arguably, when invoking such assumptions, researchers do not typically have a specific data generating process in mind that is known to satisfy the conditions; rather there is great uncertainty about the true data generating process, and the hope is that by deriving tests that are valid for a large class of data generating processes,

the true model is also covered. The primitive conditions are therefore quite possibly not a reflection of what researchers are sure is true about the data generating process, but rather an attempt to assume little in order to gain robustness. In that perspective, it seems quite natural to impose that the asymptotic rejection probability is no bigger than the nominal level for all models that satisfy $\hat{J}_T(\cdot) \rightsquigarrow W(\cdot)$. In fact, Stock (1994), White (2001, p. 179), Breitung (2002), Davidson (2002, 2007) and Müller (2008) *define* the unit root null hypothesis in terms of the convergence $T^{-1/2}u_{T, \lfloor \cdot \rfloor} \rightsquigarrow \omega W(\cdot)$, making the requirement (8) quite natural for a unit root test. \blacktriangle

On the other hand, one might argue that the only purpose of an asymptotic analysis is to generate approximations for the small sample under study. In that perspective, it is irrelevant whether interest rates are indeed mixing or not, and the only interesting question becomes whether asymptotic properties derived under an assumption of mixing are useful approximations for the small sample under study. So even in an i.i.d. setting, one might be reluctant to fully exploit all restrictions of a given model—not because it wouldn't be true that with a very large data set, a fully efficient test would be excellent, but because asymptotics might be a poor guide to the behavior of such a test in the sample under study. The robustness constraint (8) is then motivated by a concern that additional asymptotic implications of the primitive conditions beyond (2) are potentially poor approximations for the sample under study, and attempts to exploit them may lead to non-trivial size distortions.

Low Frequency Unit Root Test Example: Müller and Watson (2008b) argue that in a macroeconomic context, it makes sense to take asymptotic implications of standard models of low frequency variability seriously only over frequencies below the business cycle. So in particular, when $u_{T,t}$ is modelled as local-to-unity, then the usual asymptotic implication is the functional convergence $T^{-1/2}u_{T, \lfloor \cdot \rfloor} \rightsquigarrow \omega J_c(\cdot)$. Müller and Watson (2008b) instead derive a scale invariant (we discuss invariance in Section 3.2) point-optimal unit root test that only assumes a subset of this convergence, that is

$$\left\{ T^{-3/2} \sum_{t=1}^T \psi_l(t/T) u_{T,t} \right\}_{l=1}^q \rightsquigarrow \left\{ \omega \int_0^1 \psi_l(s) J_\theta(s) ds \right\}_{l=1}^q, \quad (9)$$

where $\psi_l(s) = \sqrt{2} \cos(\pi l s)$ and q is chosen so that the frequency of the weight functions ψ_l , $l = 1, \dots, q$ are below business cycle frequency for the span of the sample under study. The rationale is that picking q larger would implicitly imply a flat spectrum for $\Delta u_{T,t} =$

$u_{T,t} - u_{T,t-1}$ in the I(1) model over business cycle frequencies, which is not an attractive assumption for macroeconomic data. So even if one were certain that for a long enough span of observations, the functional convergence $T^{-1/2}u_{T,[\cdot T]} \rightsquigarrow \omega J_c(\cdot)$ becomes a good approximation eventually, it does not seem well-advised to exploit its implications beyond (9) for the sample under study. \blacktriangle

2.4 Main Result

In addition to the robustness constraint (8), we allow for the possibility that tests φ_T are restricted to possess further asymptotic properties. In particular, we consider

$$\liminf_{T \rightarrow \infty} \int \varphi_T dF_T(m, \theta) \geq \pi_0(\theta) \text{ for all } m \in \mathcal{M}_1, \theta \in \Theta_1 \quad (10)$$

$$\lim_{T \rightarrow \infty} \int (\varphi_T - \alpha)(f_S \circ h_T) dF_T(m, \theta) = 0 \text{ for all } m \in \mathcal{M}_0, \theta \in \bar{\Theta}_0 \text{ and } f_S \in \mathcal{F}_S. \quad (11)$$

The constraints (10) and (11) are asymptotic analogues of the constraints (6) and (7) in the limiting problem introduced above. So setting $\pi_0 = \alpha$, for instance, imposes asymptotic unbiasedness of the test φ_T in the sense that for all models $m \in \mathcal{M}_1$, the asymptotic rejection probability of φ_T under the alternative is not smaller than the nominal level. The formulation (11) of "asymptotic conditional similarity" is convenient, as it avoids explicit limits of conditional distributions; see Jansson and Moreira (2006) for discussion and references. Also, without loss of generality, we can always impose (10) and (11), since with $\pi_0 = 0$ and $\mathcal{F}_S = \{0\}$, they do not constrain the tests φ_T in any way.

The main result of this paper is that under the robustness constraint (8), efficient tests in the limiting problem φ_S^* , evaluated at sample analogues with X replaced by $X_T = h_T(Y_T)$, yield asymptotically efficient tests in the original problem involving the observations Y_T .

Theorem 1 *Let $\varphi_S^* : S \mapsto [0, 1]$ be a level- α test in the limiting problem (4) that maximizes weighted average power (5) subject to (6) and (7). Suppose φ_S^* is $P(\theta_0)$ -almost everywhere continuous for some $\theta_0 \in \Theta_0$, and define $\hat{\varphi}_T^* : \mathbb{R}^{nT} \mapsto [0, 1]$ as $\hat{\varphi}_T^* = \varphi_S^* \circ h_T$. Then*

(i) $\hat{\varphi}_T^*$ is robust, and $\lim_{T \rightarrow \infty} \text{WAP}_T(\hat{\varphi}_T^*, m) = \text{WAP}(\varphi_S^*)$ for all $m \in \mathcal{M}_1$.

(ii) For any test $\varphi_T : \mathbb{R}^{nT} \mapsto [0, 1]$ satisfying (8), (10) and (11), $\limsup_{T \rightarrow \infty} \text{WAP}_T(\hat{\varphi}_T^*, m) \leq \text{WAP}(\varphi_S^*)$ for all $m \in \mathcal{M}_1$.

Unit Root Test Example, ctd: The function $\varphi_S^* : D_{[0,1]} \mapsto [0, 1]$ is continuous at almost all realizations of W , so that part (i) of Theorem 1 shows that the test $\hat{\varphi}_T^*(Y_T) = \varphi_S^*(\hat{J}_T(\cdot)) =$

$\mathbf{1}[\exp[-\frac{1}{2}c_1(\hat{J}_T(1)^2-1)-\frac{1}{2}c_1^2 \int \hat{J}_T(s)^2 ds] > cv]$ has asymptotic null rejection probability equal to the nominal level and asymptotic weighted average power equal to $P(L(J_{c_1}) > cv)$ for all models in \mathcal{M}_0 and \mathcal{M}_1 , respectively, that is models that satisfy $\hat{J}_T(\cdot) = T^{-1/2}\hat{\omega}_T^{-1}u_{T, [\cdot, T]} \rightsquigarrow W(\cdot)$ and $\hat{J}_T(\cdot) \rightsquigarrow J_c(\cdot)$. Note that $\hat{\varphi}_T^*(Y_T)$ is asymptotically equivalent to the efficient unit root test statistic derived by ERS, so the contribution of part (i) of Theorem 1 for the unit root testing example is only to point out that $\hat{\varphi}_T^*$ has the same asymptotic properties under the null and alternative hypothesis for all models in \mathcal{M}_0 and \mathcal{M}_1 , respectively.

The more interesting finding is part (ii) of Theorem 1: For any unit root test that has higher asymptotic power than $\hat{\varphi}_T^*$ for *any* model satisfying $\hat{J}_T(\cdot) \rightsquigarrow J_{c_1}(\cdot)$, there exists a model $m \in \mathcal{M}_0$ where $\hat{J}_T(\cdot) \rightsquigarrow W(\cdot)$ for which the test has asymptotic null rejection probability larger than the nominal level. Any adaption to a non-Gaussian error distribution that leads to higher asymptotic power than ERS's test necessarily implies violation of the robustness condition (8). In other words, ERS's test is point-optimal in the class of all robust tests. \blacktriangle

The proof of part (i) of Theorem 1 follows from the definition of weak convergence, the continuous mapping theorem and dominated convergence. To gain some intuition for part (ii), consider the case where the hypotheses are simple, $\Theta_0 = \{\theta_0\}$ and $\Theta_1 = \{\theta_1\}$, and (10) and (11) do not bind. Let $L : S \mapsto \mathbb{R}$ be the Radon-Nikodym derivative of $P(\theta_1)$ with respect to $P(\theta_0)$, so that by the Neyman-Pearson Lemma, φ_S^* rejects for large values of L . For simplicity, assume that $L^i = 1/L$ is continuous and bounded. The central idea is to take the model $m \in \mathcal{M}_1$, and to 'tilt' the probabilities according to $L^i \circ h_T$ to construct a corresponding model G_T in \mathcal{M}_0 . This tilted probability distribution needs to integrate to one, so let $\kappa_T = \int (L^i \circ h_T) dF_T(m, \theta_1) = \int L^i dP_T(m, \theta_1)$, and define the measure G_T on $\mathfrak{B}(\mathbb{R}^{nT})$ via $\int_A dG_T = \kappa_T^{-1} \int_A (L^i \circ h_T) dF_T(m, \theta_1)$ for all $A \in \mathfrak{B}(\mathbb{R}^{nT})$. By construction, under G_T , the function h_T induces the measure Q_T on S , where Q_T satisfies $\int \vartheta dQ_T = \kappa_T^{-1} \int \vartheta L^i dP_T(m, \theta_1)$ for any bounded and continuous function $\vartheta : S \mapsto \mathbb{R}$. Further, the $S \mapsto \mathbb{R}$ functions ϑL^i and L^i are bounded and continuous, so that $P_T(m, \theta_1) \rightsquigarrow P(\theta_1)$ implies $\kappa_T \rightarrow \int L^i dP(\theta_1) = \int L^i L dP(\theta_0) = \int dP(\theta_0) = 1$ and $\int \vartheta L^i dP_T(m, \theta_1) \rightarrow \int \vartheta L^i dP(\theta_1) = \int \vartheta dP(\theta_0)$, so that $h_T G_T \rightsquigarrow P(\theta_0)$. Thus, by (8), $\limsup_{T \rightarrow \infty} \int \varphi_T dG_T \leq \alpha$. Furthermore, by construction, the Radon-Nikodym derivative between G_T and $F_T(m, \theta_1)$ is given by $\kappa_T(L \circ h_T)$. Therefore, by the Neyman-Pearson Lemma, the best test of $\tilde{H}_0 : Y_T \sim G_T$ against $\tilde{H}_1 : Y_T \sim F_T(m, \theta_1)$ rejects for large values of $L \circ h_T$, and no test can have a better

asymptotic level and power trade-off than this sequence of optimal tests. But $\hat{\varphi}_T^*$ also rejects for large values of $L \circ h_T$ and has the same asymptotic null rejection probability, and the result follows.

2.5 Discussion

1. Recall how the 14 papers mentioned in the introduction derive an asymptotically efficient and robust test: Initially, restrict attention to the canonical parametric version of the model of interest, usually with Gaussian i.i.d. disturbances. Call this model m^* , so that $Y_T \sim F_T(m^*, \theta)$, and for simplicity, consider the problem of testing the simple hypotheses $H_0 : \theta = \theta_0$ against $H_1 : \theta = \theta_1$. In this parametric model, $F_T(m^*, \theta_1)$ is absolutely continuous with respect to $F_T(m^*, \theta_0)$, and the small sample Likelihood Ratio statistic LR_T can be derived. The small sample optimal test in model m^* thus rejects for large values of LR_T . Express LR_T (up to asymptotically negligible terms) as a continuous function $L : S \mapsto \mathbb{R}$ of a random element $X_T^* = h_T^*(Y_T)$ that converges weakly under the null and contiguous alternative: $\text{LR}_T = L(X_T^*) + o_p(1)$, where $X_T^* \rightsquigarrow X$ with $X \sim P(\theta)$. Thus, by the continuous mapping theorem, also the likelihood ratio statistic converges weakly under the null and alternative, $\text{LR}_T \rightsquigarrow \text{LR} \sim LP(\theta)$, and the asymptotic critical value is computed from the distribution $LP(\theta_0)$. Furthermore, an asymptotically robust test statistic is given by $L(X_T)$, where $X_T = h_T(Y_T)$ is a "robustified" version of X_T^* such that $X_T \rightsquigarrow X \sim P(\theta)$ in many models m of interest, including m^* (whereas typically, $X_T^* = h_T^*(Y_T) \not\rightsquigarrow X$ for some plausible models).

Unit Root Test Example, ctd: Under i.i.d. standard normal driving disturbances, the small sample efficient unit root test rejects for large values of $\text{LR}_T = \exp[-\frac{1}{2}c_1 T^{-1}(u_{T,T}^2 - \sum_{t=1}^T (\Delta u_{T,t})^2) - \frac{1}{2}c_1^2 T^{-2} \sum_{t=1}^T u_{T,t-1}^2]$. With $h_T^*(Y_T) = X_T^* = T^{-1/2} u_{T, \lfloor \cdot \rfloor} \rightsquigarrow J_c(\cdot)$, we thus have $\text{LR}_T = L(X_T^*) + o_p(1)$ with $L(x) = \exp[-\frac{1}{2}c_1(x(1)^2 - 1) - \frac{1}{2}c_1^2 \int_0^1 x(s)^2 ds]$. The asymptotically robustified test that allows for serially correlated and non-Gaussian disturbances is based on $L(X_T)$, where $X_T = \hat{J}_T(\cdot) = T^{-1/2} \hat{\omega}_T^{-1} u_{T, \lfloor \cdot \rfloor}$. \blacktriangle

The end product of this standard approach is a test based on the statistic $L(X_T)$, with critical value computed from the distribution $LP(\theta_0)$. Now generically, this test is identical to the test $\hat{\varphi}_T^*$ of Theorem 1. This follows from a general version of LeCam's Third Lemma (see, for instance, Lemma 27 of Pollard (2001)): If the models $Y_T \sim F_T(m^*, \theta_0)$ and $Y_T \sim F_T(m^*, \theta_1)$ are contiguous with likelihood ratio statistic LR_T , and under $Y_T \sim F_T(m^*, \theta_0)$,

$(LR_T, X_T^*) \rightsquigarrow (L(X), X)$ with $X \sim P(\theta_0)$ and some function $L : S \mapsto \mathbb{R}$, then under $Y_T \sim F_T(m^*, \theta_1)$, $X_T^* \rightsquigarrow X \sim Q$, and the Radon-Nikodym derivative of Q with respect to $P(\theta_0)$ is equal to L . So if it is known that under $Y_T \sim F_T(m^*, \theta_1)$, $X_T^* \rightsquigarrow X \sim P(\theta_1)$, then it must be the case that $Q = P(\theta_1)$, and $L(X)$ is recognized as the Neyman-Pearson test statistic of the limiting problem $H_0^{lp} : \theta = \theta_0$ against $H_1^{lp} : \theta = \theta_1$ with $X \sim P(\theta)$ observed. The test that rejects for large values of $L(X_T)$ is thus simply the efficient test of this limiting problem, evaluated at sample analogues. This explains why in the unit root example the test $\hat{\varphi}_T^*$ of Theorem 1 had to be asymptotically equivalent to ERS's statistic.

2. This standard construction of tests starting from the canonical parametric model m^* , i.e. rejecting for large values of $L(X_T)$, is by construction asymptotically efficient in model m^* , and by part (i) of Theorem 1, it has the same asymptotic local power for all models $m \in \mathcal{M}_1$ for which $X_T \rightsquigarrow X \sim P(\theta_1)$. This does not, however, make the test $L(X_T)$ necessarily overall asymptotically efficient: It might be that there exists another test with the same asymptotic power in model m^* , and higher asymptotic power for at least some other models m for which $X_T \rightsquigarrow X \sim P(\theta_1)$. The semiparametrically efficient unit root test by Jansson (2008) is an example of a test with the same asymptotic power as ERS's test for Gaussian i.i.d. disturbances, and higher asymptotic power for some other driving disturbances.

Now part (ii) of Theorem 1 shows that whenever a test has higher asymptotic power than $\hat{\varphi}_T^*$ for some alternative model $m \in \mathcal{M}_1$, then it cannot satisfy the robustness constraint (8). So any partial adaption to models $m \neq m^*$, if successful, necessarily implies the existence of a model $m \in \mathcal{M}_0$ for which the test has asymptotic rejection probability larger than the nominal level. So in particular, Theorem 1 implies the existence of a double array process $(u_{T,1}, \dots, u_{T,T})'$ satisfying $T^{-1/2} \hat{\omega}_T^{-1} u_{T, [T]} \rightsquigarrow W(\cdot)$ for which Jansson's (2008) test has asymptotic rejection probability larger than the nominal level.

In other words, under the robustness constraint (8), Theorem 1 shows $\hat{\varphi}_T^*$ to be an overall asymptotically efficient test, because no test can exist with higher asymptotic (weighted average) power for *any* model $m \in \mathcal{M}_1$.

3. In this sense, Theorem 1 implies a particular version of an asymptotic essentially complete class result for the hypothesis test (1): Set π_0 in (10) equal to the power function of an admissible test in the limiting problem (4), so that (10) effectively determines φ_S^* . The theorem then shows that no test φ_T of (1) can be robust in the sense of (8) and have

higher asymptotic power than $\hat{\varphi}_T^*$ uniformly over $\theta \in \Theta_1$. As long as all admissible φ_S^* are $P(\theta_0)$ -almost everywhere continuous, the resulting tests $\hat{\varphi}_T^*$ thus form an "essentially complete class of asymptotically admissible robust tests" of the original problem (1).

In particular, if a uniformly most powerful test exists in the limiting problem, then φ_S^* is this test for any weighting function w , and repeated application of Theorem 1 with w having point mass for any $\theta \in \Theta_1$ then shows that the test $\hat{\varphi}_T^*$ is correspondingly asymptotically uniformly most powerful.

4. When Θ is of finite dimension, the models of interest \mathcal{M}_0 and \mathcal{M}_1 may be called semiparametric, with the weak convergence assumption constituting the nonparametric aspect. This raises the question how Theorem 1 relates to the more traditional concept of semiparametrically efficient inference in models with locally asymptotically normal (LAN) likelihoods, which has seen numerous recent applications in econometrics (see, for instance Bickel, Klaassen, Ritov, and Wellner (1998), Newey (1990) or chapter 25 of van der Vaart (1998) for an introduction).

For expositional ease, consider T scalar observations $Y_T = (y_{T,1}, \dots, y_{T,T})'$ with mean $E[y_{T,t}] = \mu$ and variance $E[(y_{T,t} - \mu)^2] = 1$, and the hypotheses are $H_0 : \mu = \mu_0$ against $H_1 : \mu > \mu_0$. Reparametrize $\theta = \sqrt{T}(\mu - \mu_0)$, so that in terms of θ the problem becomes $H_0 : \theta = 0$ against $H_1 : \theta > 0$. It is natural to take the sample mean $\bar{y}_T = T^{-1} \sum_{t=1}^T y_{T,t}$ as the basis for inference, so consider the monotone transformation $X_T = h_T(Y_T) = \sqrt{T}(\bar{y}_T - \mu_0) \rightsquigarrow X \sim \mathcal{N}(\theta, 1)$, where the convergence stems from a Central Limit Theorem. The likelihood ratio statistic of $X \sim \mathcal{N}(\theta_1, 1)$ against $X \sim \mathcal{N}(0, 1)$ equals $L(X) = \exp[\theta_1 X - \frac{1}{2}\theta_1^2]$, so that by the Neyman-Pearson Lemma, $\varphi_S^*(X) = \mathbf{1}[X > 1.645]$ is the the uniformly most powerful 5% level test in the limiting problem of observing $X \sim \mathcal{N}(\theta, 1)$. Thus, by Theorem 1, $\hat{\varphi}_T^*(Y_T) = \varphi_S^*(X_T)$ is the best 5% level test among all tests whose asymptotic null rejection probability is no bigger than 5% whenever $X_T \rightsquigarrow \mathcal{N}(0, 1)$.

Ignoring the unboundedness of $L^i = 1/L$ for simplicity, the basic argument in the proof of Theorem 1 starts from any alternative model $F_T(m, \theta_1)$ (so that under $F_T(m, \theta_1)$, $X_T \sim \mathcal{N}(\theta_1, 1)$) and constructs a corresponding null model G_T by tilting the probabilities according to $L^{-1} \circ h_T = \exp[-\theta_1 X_T + \frac{1}{2}\theta_1^2]$, that is $\int_A dG_T = \kappa_T^{-1} \int_A (L^i \circ h_T) dF_T(m, \theta_1)$ for all $A \in \mathfrak{B}(\mathbb{R}^T)$, where $\kappa_T \rightarrow 1$ insures that $\int_{\mathbb{R}^T} dG_T = 1$. If under the alternative model $F_T(m, \theta_1)$,

$\{y_{T,t}\}_{t=1}^T$ is i.i.d. with log density $y \mapsto g(y - T^{-1/2}\theta_1)$, we obtain

$$\int_A dG_T = \int_A \prod_{t=1}^T \exp[g(y_{T,t} - T^{-1/2}\theta_1) - T^{-1/2}\theta_1 y_{T,t} + \frac{1}{2}T^{-1}\theta_1^2 - T^{-1} \ln \kappa_T] dY_T \quad (12)$$

so that the tilted model G_T is also i.i.d., with log density $g(y - T^{-1/2}\theta_1) - T^{-1/2}\theta_1 y + \frac{1}{2}T^{-1}\theta_1^2 - T^{-1} \ln \kappa_T$.

Now consider the analysis of inference about the mean in an i.i.d. scalar series using traditional semiparametric techniques. Write ε for a mean zero random variable with distribution $y_{T,t} - \mu$, and embed the log density g of ε in a smooth parametric model g_η with parameter $\eta \in \mathbb{R}$. Under the local parametrization $(\theta, v) = \sqrt{T}(\mu - \mu_0, \eta - \eta_0)$ and some regularity conditions, we obtain that the log-likelihood process satisfies

$$\sum_{t=1}^T (g_{\eta_0 + T^{-1/2}v}(y_{T,t} - T^{-1/2}\theta) - g_{\eta_0}(y_{T,t})) = S'_T \begin{pmatrix} \theta \\ v \end{pmatrix} - \frac{1}{2} \begin{pmatrix} \theta \\ v \end{pmatrix}' \mathcal{I} \begin{pmatrix} \theta \\ v \end{pmatrix} + o_p(1) \quad (13)$$

where $S_T = T^{-1/2}(\sum_{t=1}^T \ell_\mu(y_{T,t}), \sum_{t=1}^T \ell_\eta(y_{T,t}))$, $\mathcal{I} = ((\mathcal{I}_{\mu\mu}, \mathcal{I}_{\mu\eta})', (\mathcal{I}_{\mu\eta}, \mathcal{I}_{\eta\eta})')$, ℓ_μ and ℓ_η are $\mathbb{R} \mapsto \mathbb{R}$ functions with $E[\ell_\mu(\varepsilon)^2] = \mathcal{I}_{\mu\mu}$, $E[\ell_\eta(\varepsilon)^2] = \mathcal{I}_{\eta\eta}$, $E[\ell_\mu(\varepsilon)\ell_\eta(\varepsilon)] = \mathcal{I}_{\mu\eta}$ and $E[\varepsilon(\ell_\mu(\varepsilon), \ell_\eta(\varepsilon))] = (1, 0)$. It is useful to think of ℓ_μ and ℓ_η as appropriate generalizations of the usual score functions $\partial g_\eta(y - \mu)/\partial \mu|_{\mu=\mu_0}$ and $\partial g_\eta(y)/\partial \eta|_{\eta=\eta_0}$, respectively, so that under weak regularity conditions, $E[\varepsilon(\ell_\mu(\varepsilon), \ell_\eta(\varepsilon))] = (1, 0)$ follows from the identity $\int y \exp g_\eta(y - \mu) dy = \mu$. Furthermore, by LeCam's Third Lemma, $S_T \rightsquigarrow \mathcal{N}(\mathcal{I}(\theta, v)', \mathcal{I})$ when the true log density of $y_{T,t}$ is given by $g_{\eta_0 + T^{-1/2}v}(y - T^{-1/2}\theta)$. For any $\theta_1 > 0$, consider the best test of the simple null hypothesis $H_0 : (\theta, v) = (0, v_0)$ with $v_0 = \mathcal{I}_{\eta\eta}^{-1} \mathcal{I}_{\mu\eta} \theta_1$ against the single alternative $H_1 : (\theta, v) = (\theta_1, 0)$. In large samples, the Neyman-Pearson test rejects for large values of $((\theta_1, -v_0) \mathcal{I}(\theta_1, -v_0)')^{-1/2} (\theta_1, -v_0) S_T \rightsquigarrow \mathcal{N}(a, 1)$, where $a = 0$ under the null hypothesis, and $a = \theta_1 \sqrt{\mathcal{I}_{\mu\mu} - \mathcal{I}_{\mu\eta}^2 / \mathcal{I}_{\eta\eta}}$ under the alternative. Because $a = 0$ for any local parameter of the form $(0, v)$, the choice of $v_0 = \mathcal{I}_{\eta\eta}^{-1} \mathcal{I}_{\mu\eta} \theta_1$ is least favorable, so that by Theorem 3.8.1 of Lehmann and Romano (2005), this test is the best test against H_1 also under the composite null hypothesis that leaves v unrestricted. Furthermore, this test does not depend on θ_1 , and is thus uniformly most powerful against alternatives of the form $(\theta, 0)$ with $\theta > 0$.

Now Stein's (1956) key insight underlying semiparametric efficiency calculations is that inference with g unknown is at least as hard as inference in any of the smooth parametric submodels with parameter η . Because the power of the best test in any particular submodel

increases in $\mathcal{I}_{\mu\mu} - \mathcal{I}_{\mu\eta}^2/\mathcal{I}_{\eta\eta}$, the least favorable submodel is the one that minimizes this quantity, subject to the constraint that $E[\varepsilon(\ell_\mu(\varepsilon), \ell_\eta(\varepsilon))] = (1, 0)$. This is solved by $\ell_\eta(\varepsilon) = \ell_\mu(\varepsilon) - \varepsilon$, for which $\mathcal{I}_{\eta\eta} = \mathcal{I}_{\mu\eta} = \mathcal{I}_{\mu\mu} - 1$ and $a = \theta_1$ under the alternative. But $X_T = \sqrt{T}(\bar{y}_T - \mu_0)$ achieves this non-centrality parameter, so that $\hat{\varphi}_T^*(Y_T) = \mathbf{1}[X_T > 1.645]$ is also the semiparametrically efficient 5% level test.

As just discussed, the rationale for this semiparametric efficiency claim is that for any alternative model with log density $g(y - T^{-1/2}\theta_1)$, there exists (or, more precisely, one can arbitrarily well approximate) a least favorable smooth parametric submodel g_η with $g = g_{\eta_0}$ and $\ell_\eta(\varepsilon) = \ell_\mu(\varepsilon) - \varepsilon$, and a least favorable null value $\nu_0 = \mathcal{I}_{\eta\eta}^{-1}\mathcal{I}_{\mu\eta}\theta_1 = \theta_1$ of the parameter $\eta = \eta_0 + T^{-1/2}\nu_0$ so that by (13), the likelihood ratio process satisfies $\sum_{t=1}^T (g_{\eta_0}(y_{T,t} - T^{-1/2}\theta_1) - g_{\eta_0+T^{-1/2}\theta_1}(y_{T,t})) = \theta_1 T^{-1/2} \sum_{t=1}^T y_{T,t} - \frac{1}{2}\theta_1^2 + o_p(1)$. Solving this for the least favorable null log density, we obtain

$$\sum_{t=1}^T g_{\eta_0+T^{-1/2}\theta_1}(y_{T,t}) = \sum_{t=1}^T (g_{\eta_0}(y_{T,t} - T^{-1/2}\theta_1) - \theta_1 T^{-1/2} y_{T,t} + \frac{1}{2} T^{-1} \theta_1^2) + o_p(1),$$

which is essentially the same as the tilted model (12). So not only do both approaches to the question of efficient tests lead to the same test $\hat{\varphi}_T^*(Y_T) = \mathbf{1}[X_T > 1.645]$, but also the least favorable null model that 'defeats' attempts of learning from the data about μ through statistics other than X_T is (at least approximately) the same in both approaches.

This basic logic and equivalence result extend to general semiparametric LAN models, where semiparametrically efficient tests amount to efficient tests in a pseudo model that treats the semiparametric efficient estimator as a single normal observation with known covariance matrix—see Choi, Hall, and Schick (1996) and Chapter 25.6 of van der Vaart (1998) for a general and rigorous discussion of the semiparametric approach to testing outlined above. Quite generally, then, $\hat{\varphi}_T^*$ of Theorem 1 equals the usual semiparametric efficient test when $X_T = h_T(Y_T)$ is a suitably scaled function of a semiparametric efficient estimator in a LAN model, and the underlying logic for these efficiency claims are very similar. Despite the different starting point and underlying regularity assumptions, the approach of this paper does not yield any different conclusion in such models. The semiparametric efficiency result is arguably sharper, though, as it accommodates moment conditions and, say, the i.i.d. property of the observations, whereas Theorem 1 makes a less specific efficiency claim.

5. One advantage of Theorem 1 relative to the semiparametric efficiency approach is the general nature of its claim. The result requires only minimal regularity conditions

beyond the basic weak convergence assumption, and covers a wide range of non-standard testing problems. In contrast, relatively little is known about traditional semiparametric efficiency in models outside the LAN class, and deriving such results is non-trivial and involves additional regularity assumptions (cf. Jansson (2008)).

When Theorem 1 yields a less powerful test than what is obtained from a traditional semiparametric efficiency calculation, one might try to reconcile the results by either assuming additional weak convergences (which potentially increases the power of $\hat{\varphi}_S^*$) or by considering less constrained semiparametric models (which potentially decreases the power of the semiparametrically efficient test). For the latter, the inspection of the form of the tilted model G_T in the proof Theorem 1 may suggest an appropriate relaxation of the constraints.

Unit Root Test Example, ctd. Proceeding as in Rothenberg and Stock (1997) and Jansson (2008) yields that the log of the likelihood ratio statistic for testing $H_0 : c = 0$ against $H_1 : c = c_1$ in an AR(1) with coefficient $1 - c/T$ and i.i.d. driving errors with the same distribution as ε satisfies $-c_1 T^{-1} \sum_{t=1}^T \ell_\mu(\Delta u_{T,t}) u_{T,t-1} - \frac{1}{2} c_1^2 \mathcal{I}_{\mu\mu} T^{-2} \sum_{t=1}^T u_{T,t-1}^2 + o_p(1)$, with ε , ℓ_μ and $\mathcal{I}_{\mu\mu}$ as defined in comment 4 above. Heuristically, the tilting by $1/L$ is approximately equal to a tilting by $1/\text{LR}_T$, and a model tilted by $1/\text{LR}_T$ has joint log-density $\sum_{t=1}^T [g(\Delta u_{T,t}) - c_1 T^{-1} (\ell_\mu(\Delta u_{T,t}) - \Delta u_{T,t}) u_{T,t-1} - \frac{1}{2} c_1^2 T^{-2} (\mathcal{I}_{\mu\mu} - 1) u_{T,t-1}^2] + o_p(1)$. This is recognized as an expansion of $\sum_{t=1}^T g_{\eta_0 - c_1/T}^M(\Delta u_{T,t}, u_{T,t-1})$ with $\exp(g_\eta^M(y, x)) = e^{g(y)}(1 + (\eta - \eta_0)(\ell_\mu(y) - y)x)$ for $\eta \in \mathbb{R}$. Note that $\int \exp(g_\eta^M(y, x)) dy = 1$ and $\int y \exp(g_\eta^M(y, x)) dy = 0$ for all x and η , suggesting that $\sum_{t=1}^T g_{\eta_0 - c_1/T}^M(\Delta u_{T,t}, u_{T,t-1})$ are the log Markov kernels of a time homogeneous first order Markov martingale for $\{u_{T,t}\}_{t=1}^T$. As $(\ell_\mu(y) - y)x$ is unbounded, though, $g_\eta^M(y, x)$ is not well defined even in the T^{-1} neighborhood of η_0 . But one could presumably employ an appropriate approximation scheme, so that an analysis analogous to Jansson (2008) would yield $\hat{\varphi}_T^*$ to be the semiparametrically point-optimal unit root tests in the first order time homogeneous Markov model against AR(1) alternatives. A detailed analysis of this conjecture is beyond the scope of this paper. \blacktriangle

6. Weak convergence statements of the form (2) and (3) can be viewed as a way of expressing regularity one is willing to impose on some inference problem. Implicitly, this is standard practice: invoking standard normal asymptotics for the OLS estimator of the largest autoregressive root ρ is formally justified for any value of $|\rho| < 1$, but effectively amounts to the assumption that the true parameter in the sample under study is not close

to the local-to-unity region. Similarly, a choice of weak vs strong instrument asymptotics or local vs non-local time varying parameter asymptotics expresses knowledge of regularity in terms of weak convergences.

In some instances, it might be natural to express all regularity that one is willing to impose in this form, and Theorem 1 then shows that $\hat{\varphi}_T^*$ efficiently exploits this information. The i.i.d. nature of (standard) cross sectional data is not easily embedded in a weak convergence statement, so that such a starting point is much more convincing in a time series context. Also, interesting high level weak convergence assumptions are certainly not entirely arbitrary, but derive their plausibility from the knowledge that there exists a range of underlying primitive conditions that would imply them.

If one expresses regularity of a problem in terms of weak convergences, one faces a choice of what to assume. Of course, additional weak convergence assumptions cannot reduce the information about θ in the limiting problem. But not all weak convergences are relevant for deciding between H_0 and H_1 : Whenever the conditional distribution of the additional limiting element does not depend on θ , the efficient test remains unaltered. This holds, for example, for any additional convergence in probability to a constant limiting element (whose value does not depend on θ). In the context of unit root tests, for instance, ERS's test remains asymptotically efficient in the sense of Theorem 1 if in addition to $\hat{J}_T(\cdot) \rightsquigarrow J_c(\cdot)$, the average sample skewness of $\Delta u_{T,t} = u_{T,t} - u_{T,t-1}$, $T^{-1} \sum_{t=1}^T (\Delta u_{T,t})^3$, is assumed to converge to zero in probability, or that $T^{-1/2} \sum_{t=1}^{\lfloor \cdot T \rfloor} (\Delta u_{T,t})^3 \rightsquigarrow W_\Delta(\cdot)$ for any $c \geq 0$, where $W_\Delta(\cdot)$ is a Wiener process independent of J_c . Additional weak convergence restrictions of this type strengthen the efficiency claim by weakening the robustness requirement (8) of tests, since \mathcal{M}_0 becomes a smaller set, but, of course, the resulting efficiency correspondingly only holds for the smaller alternative set \mathcal{M}_1 .

At the same time, one might also be reluctant to impose the full extent of the 'usual' weak convergence assumption, and in general, this leads to less powerful inference. The efficiency claim of Theorem 1 then shows that it is impossible to use data-dependent methods to improve inference for more regular data while still remaining robust in the sense of (8).

Low Frequency Unit Root Test Example, ctd: Since (9) is strictly weaker than the standard assumption $T^{-1/2} u_{T, \lfloor \cdot T \rfloor} \rightsquigarrow \omega J_c(\cdot)$, Müller and Watson's (2008) low-frequency unit root test is less powerful than ERS's test. It is nevertheless point-optimal in the sense of efficiently extracting all regularity contained in the weaker statement (9): Theorem 1 implies

that it is impossible to let the data decide whether (9) holds for q larger than assumed (that is, whether the local-to-unity model provides good approximations also over business cycle frequencies), and to conduct more powerful inference if it is, without inducing size distortions for some model satisfying (9) for $c = 0$. \blacktriangle

7. No matter how one views the desirability of the robustness constraint (8), one appeal of Theorem 1 is that it suggests a general method for constructing reasonable tests. In nonlinear and/or dynamic models, it might be difficult to derive the small sample likelihood ratio statistic, even under strong parametric assumptions, while high level weak convergence properties might be easier to think about. The problem of testing for parameter instability in a general GMM framework, as considered by Sowell (1996) and discussed in detail in Section 4.3 below, or the weak instrument problem in a general GMM framework, as considered by Stock and Wright (2000), arguably fall into this class. For such problems, efficient tests of the limiting problem, evaluated at sample analogues, are a natural starting point for sensible tests in the original problem.

3 Extensions

3.1 Consistently Estimable Nuisance Parameters

Suppose the testing problem (1) involves an additional nuisance parameter $\gamma \in \Gamma$, where Γ is a metric space, so that now $Y_T \sim F_T(m, \theta, \gamma)$, and the null and alternative hypotheses become

$$H_0 : (\theta, \gamma) \in (\Theta_0, \Gamma) \quad \text{against} \quad H_1 : (\theta, \gamma) \in (\Theta_1, \Gamma). \quad (14)$$

Suppose γ can be consistently estimated by the estimator $\hat{\gamma}_T$ under the null and alternative hypothesis. For a fixed value of $\gamma = \gamma_0$, that is when Γ is a singleton, this is a special case of what is covered by Theorem 1, as discussed in comment 6. But when Γ is not a singleton, the analysis above is not immediately applicable, because the limiting measures of X were assumed to be mutually absolutely continuous for all parameter values.

Thus denote by $P_T^e(m, \theta, \gamma)$ ('e' for extended) the distribution of $(\hat{\gamma}_T, X_T) = h_T^e(Y_T)$ when $Y_T \sim F_T(m, \theta, \gamma)$, so that the weak convergence assumption analogous to (2) and (3) now becomes

$$P_T^e(m, \theta, \gamma) \rightsquigarrow P^e(\theta, \gamma) \quad \text{pointwise for all } (\theta, \gamma) \in (\Theta \times \Gamma) \quad (15)$$

where $P^e(\theta, \gamma)$ is the product measure between the measure $P(\theta, \gamma)$ of X (which might depend on γ) on $\mathfrak{B}(S)$ and the degenerate probability measure on $\mathfrak{B}(\Gamma)$ that puts all mass on the point γ . The limiting problem now becomes testing (14) with $X \sim P(\theta, \gamma)$ observed and γ known, and optimal tests φ_S^{e*} in the limiting problem are indexed by $\gamma \in \Gamma$, $\varphi_S^{e*} : \Gamma \times S \mapsto [0, 1]$, so that for each $\gamma_0 \in \Gamma$, $\varphi_S^{e*}(\gamma_0, \cdot)$ is the weighted average power maximizing test of (14) for $\gamma = \gamma_0$, possibly subject to constraints of the form (6) and (7).

Now as long as the test φ_S^{e*} is, for each $\gamma \in \Gamma$, $P^e(\gamma, \theta_0)$ -almost everywhere continuous, the same arguments as employed in the proof of Theorem 1 part (i) imply that under (15) and for each $\gamma \in \Gamma$, the test $\hat{\varphi}_T^{e*} = \varphi_S^{e*}(\hat{\gamma}_T, X_T)$ has the same asymptotic size and power as φ_S^{e*} does. Furthermore, for any value of $\gamma = \gamma_0$, one can invoke Theorem 1 (ii) as above to conclude that no test $\hat{\varphi}_T$ can exist with better asymptotic weighted average power, for any model satisfying $P_T^e(m, \theta, \gamma_0) \rightsquigarrow P^e(\theta, \gamma_0)$. Problems that involve consistently estimable parameters γ with an impact on the efficient limiting tests are thus covered by the results of Theorem 1 under an additional assumption of the family of efficient limiting tests, indexed by γ , to depend on γ sufficiently smoothly.

Unit Root Test Example, ctd: Instead of $T^{-1/2}\hat{\omega}_T^{-1}u_{T, \lfloor \cdot T \rfloor} = \hat{J}_T(\cdot) \rightsquigarrow J_c(\cdot)$, consider the weak convergences $h_T^e(Y_T) = (\hat{\omega}_T^2, T^{-1/2}u_{T, \lfloor \cdot T \rfloor}) \rightsquigarrow (\omega^2, \omega J_c(\cdot))$ as a starting point. In the limiting problem, $X = \omega J_c(\cdot)$ is observed with ω^2 known, and the point-optimal test is of the form $\varphi_S^{e*}(\omega^2, X) = \mathbf{1}[\exp[-\frac{1}{2}c_1(\omega^{-2}X(1)^2 - 1) - \frac{1}{2}c_1^2\omega^{-2} \int_0^1 X(s)^2 ds] > cv]$. A calculation shows this to be a continuous function $(0, \infty) \times D_{[0,1]} \mapsto \mathbb{R}$ for almost all realizations of (ω^2, J_0) , so the test $\varphi_S^{e*}(\hat{\omega}_T^2, T^{-1/2}u_{T, \lfloor \cdot T \rfloor})$ is asymptotically efficient among all unit root tests with correct asymptotic null rejection probability whenever $(\hat{\omega}_T^2, T^{-1/2}u_{T, \lfloor \cdot T \rfloor}) \rightsquigarrow (\omega^2, \omega J_0(\cdot))$ against all models satisfying $(\hat{\omega}_T^2, T^{-1/2}u_{T, \lfloor \cdot T \rfloor}) \rightsquigarrow (\omega^2, \omega J_{c_1}(\cdot))$. \blacktriangle

3.2 Invariance

The majority of efficient tests for nonstandard problems cited in the introduction rely on invariance considerations. In the framework here, invariance may be invoked at two levels: On the one hand, one might consider a weak convergence as a starting point that is a function of a small sample maximal invariant. On the other hand, invariance might instead be employed in the limiting problem as a way of dealing with nuisance parameters. This subsection discusses the link between these two notions, and the interaction of the concept of invariance with the efficiency statements of Theorem 1.

The first case is entirely straightforward: suppose $\phi_T(Y_T)$ with $\phi_T : \mathbb{R}^{nT} \mapsto \mathbb{R}^{nT}$ is a maximal invariant to some group of transformations. By Theorem 6.2.1 of Lehmann and Romano (2005), all invariant tests can be written as functions of a maximal invariant. So if h_T of Section 2.1 is of the form $h_T = h_T^\phi \circ \phi_T$, then Theorem 1 applies and yields an asymptotic efficiency statement among all invariant tests relative to a robustness constraint (8) in terms of the weak convergence of a function of the small sample maximal invariant.

Unit Root Test Example, ctd: Consider the problem of testing for a unit root in a model with unknown mean, and suppose $\omega = 1$ is known for simplicity. A maximal invariant is given by the demeaned data $\{\hat{u}_{T,t}\}_{t=1}^T$ with $\hat{u}_{T,t} = y_{T,t} - \bar{y}_T$, where $\bar{y}_T = T^{-1} \sum_{t=1}^T y_{T,t}$ and $Y_T = (y_{T,1}, \dots, y_{T,T})'$. The typical model satisfies $T^{-1/2} \hat{u}_{T, \lfloor \cdot T \rfloor} \rightsquigarrow J_c^\mu(\cdot)$, where $J_c^\mu(s) = J_c(s) - \int_0^1 J_c(l) dl$. Theorem 1 now shows that rejecting for large values of $L^\mu(T^{-1/2} \hat{u}_{T, \lfloor \cdot T \rfloor})$ is asymptotically point-optimal among all tests whose asymptotic null rejection probability is at most α whenever $T^{-1/2} \hat{u}_{T, \lfloor \cdot T \rfloor} \rightsquigarrow J_c^\mu(\cdot)$, where L^μ is the Radon-Nikodym derivative of the probability measure of $J_{c_1}^\mu$ with respect to the measure of J_0^μ .³ Rejecting for large values of $L^\mu(T^{-1/2} \hat{u}_{T, \lfloor \cdot T \rfloor})$ is therefore the asymptotically point-optimal translation invariant test among all tests that do not overreject asymptotically whenever $T^{-1/2} \hat{u}_{T, \lfloor \cdot T \rfloor} \rightsquigarrow J_0^\mu(\cdot)$. \blacktriangle

In the second case, one considers the typical weak convergence in a model with nuisance parameters, and applies invariance only in the limiting problem. Formally, let $\tilde{g} : R \times S \mapsto S$ be such that the $S \mapsto S$ functions $x \mapsto \tilde{g}(r, x)$, indexed by $r \in R$, form the group $\tilde{\mathcal{G}}$. Further, suppose $\tilde{\phi} : S \mapsto S$ is a maximal invariant to $\tilde{\mathcal{G}}$, so that the efficient invariant test in the limiting problem $\varphi_S^{\phi^*}$ is the efficient test of $H_0^{\phi, lp} : \Xi \sim \tilde{\phi}P(\theta)$, $\theta \in \Theta_0$ against $H_1^{\phi, lp} : \Xi \sim \tilde{\phi}P(\theta)$, $\theta \in \Theta_1$, where $\Xi = \tilde{\phi}(X)$. It is not clear whether or in which sense this test, evaluated at sample analogues would be asymptotically efficient.

Unit Root Test Example, ctd: In the parametrization $y_{T,t} = u_{T,t} + T^{1/2} \alpha_y$, we obtain with $\theta = (c, \alpha_y)'$ that $T^{-1/2} y_{T, \lfloor \cdot T \rfloor} \rightsquigarrow J_\theta^y(\cdot)$, where $J_\theta^y(s) = J_c(s) + \alpha_y$. The parameter α_y is a nuisance parameters in the limiting problem. Define the transformations $\tilde{g} : \mathbb{R} \times D_{[0,1]} \mapsto D_{[0,1]}$ as $\tilde{g}(r, x) = x(\cdot) + r$, with $r \in R = \mathbb{R}$. The limiting problem is seen to be invariant to these transformations, and $\tilde{\phi} : D_{[0,1]} \mapsto D_{[0,1]}$ with $\tilde{\phi}(x) = x(\cdot) - \int_0^1 x(l) dl$, is

³Under the assumption of $u_{T,0} = 0$ for the initial condition, $L^\mu(x) = L(x(\cdot) - x(0))$, where L is the Radon-Nikodym derivative of the probability measure of J_{θ_1} with respect to the measure of J_0 , so that rejecting for large values of $L^\mu(T^{-1/2} \hat{u}_{T, \lfloor \cdot T \rfloor})$ leads to the same asymptotic power as without translation invariance. This equivalence does not hold, however, when the initial condition $u_{T,0}$ is of the same order of magnitude as $u_{T, \lfloor sT \rfloor}$ for $s > 0$. See Müller and Elliott (2003) for discussion.

a maximal invariant. Since $\tilde{\phi}(J_\theta^y) \sim J_c^\mu(\cdot)$, the point-optimal invariant test in the limiting problem rejects for large values of $L^\mu(J_\theta^y)$. This test, evaluated at sample analogues, yields $L^\mu(T^{-1/2}\hat{u}_{T,[.T]})$, just as above.

Even though in this example, the efficient invariant test in the limiting problem, evaluated at sample analogues, is small sample invariant, one still cannot claim this test to be the asymptotically point-optimal test among all small sample invariant tests with nominal asymptotic null rejection probability whenever $T^{-1/2}y_{T,[.T]} \rightsquigarrow J_0(\cdot) + \alpha_y$ with $\alpha_y \in \mathbb{R}$. The reason is that the set of models satisfying $T^{-1/2}y_{T,[.T]} \rightsquigarrow J_c(\cdot) + \alpha_y$ is a proper subset of the set of models satisfying $T^{-1/2}\hat{u}_{T,[.T]} \rightsquigarrow J_\theta^\mu(\cdot)$. The efficiency of $L^\mu(T^{-1/2}\hat{u}_{T,[.T]})$ noted above is thus relative to a more stringent robustness constraint, and it remains unclear whether an efficiency claim can also be made relative to the weaker constraint of asymptotic size control whenever $T^{-1/2}y_{T,[.T]} \rightsquigarrow J_0(\cdot) + \alpha_y$. \blacktriangle

We now show how one can make an asymptotic efficiency claim when invariance is employed in the limiting problem by relating the limiting group of transformations to a sequence of small sample groups. So for each T , suppose the measurable function $g_T : R \times \mathbb{R}^{nT} \mapsto \mathbb{R}^{nT}$ is such that the $\mathbb{R}^{nT} \mapsto \mathbb{R}^{nT}$ functions $y \mapsto g_T(r, y)$, indexed by $r \in R$, form the group \mathcal{G}_T . Let $\phi_T : \mathbb{R}^{nT} \mapsto \mathbb{R}^{nT}$ be a maximal invariant of \mathcal{G}_T . Assume that the small sample and limiting invariance correspond in the sense that the small sample maximal invariant converges weakly to the maximal invariant of the limiting problem, i.e.

$$(h_T \circ \phi_T)F_T(m, \theta) \rightsquigarrow \tilde{\phi}P(\theta) \quad \text{pointwise for } \theta \in \Theta_0 \quad (16)$$

$$(h_T \circ \phi_T)F_T(m, \theta) \rightsquigarrow \tilde{\phi}P(\theta) \quad \text{pointwise for } \theta \in \Theta_1. \quad (17)$$

Let \mathcal{M}_0^ϕ and \mathcal{M}_1^ϕ the set of models m satisfying (16) and (17), respectively. Since $\varphi_S^{\phi*}$ was assumed to be the efficient test of $H_0^{\phi,lp} : \Xi \sim \tilde{\phi}P(\theta), \theta \in \Theta_0$ against $H_1^{\phi,lp} : \Xi \sim \tilde{\phi}P(\theta), \theta \in \Theta_1$, one can apply Theorem 1 to conclude that $\varphi_S^{\phi*} \circ h_T \circ \phi_T$ is the asymptotically efficient small sample invariant (with respect to \mathcal{G}_T) robust test *relative to the weak convergence assumptions described by \mathcal{M}_0^ϕ and \mathcal{M}_1^ϕ* . As noted in the unit root example above, however, one cannot conclude that $\varphi_S^{\phi*} \circ h_T \circ \phi_T$ is also the asymptotically efficient small sample invariant robust test relative to the weak convergence $h_T F_T(m, \theta) \rightsquigarrow P(\theta)$, i.e. relative to the (typically strictly smaller) sets \mathcal{M}_0 and \mathcal{M}_1 .

It would be possible to draw this additional conclusion if for small sample invariant tests, the constraints (6), (7) and (8) with respect to models in \mathcal{M}_j are implied by the analogous

constraints with respect to \mathcal{M}_j^ϕ , $j = 0, 1$: For each model $m \in \mathcal{M}_j^\phi$, there must exist a corresponding model in \mathcal{M}_j that is identical up to a transformation in \mathcal{G}_T . The following Theorem provides conditions under which this is the case.

Theorem 2 *Suppose (i) R is a separable and complete metric space; (ii) the mapping $x \mapsto \tilde{g}(r, x)$ is continuous for all $r \in R$; (iii) there exists a measurable function $\tilde{\rho} : S \mapsto R$ such that $x = \tilde{g}(\tilde{\rho}(x), \tilde{\phi}(x))$ for all $x \in S$; (iv) $\sup_{r \in R, y \in \mathbb{R}^{nT}} d_S(h_T(g_T(r, y)), \tilde{g}(r, h_T(y))) \rightarrow 0$, where d_S is the metric on S ; (v) $\tilde{\phi}$ and ϕ_T select specific orbits, i.e. for all $y \in \mathbb{R}^{nT}$ and $x \in S$, there exist $r_y, r_x \in R$ so that $\phi_T(y) = g_T(r_y, y)$ and $\tilde{\phi}(x) = \tilde{g}(r_x, x)$. Then for any $m \in \mathcal{M}_j^\phi$ and $\theta \in \Theta_j$, $j = 0, 1$, there exists a sequence of measures G_T on $\mathfrak{B}(\mathbb{R}^{nT})$ such that $\phi_T F_T(m, \theta) = \phi_T G_T$ and $h_T G_T \rightsquigarrow P(\theta)$.*

Unit Root Test Example, ctd: With $g_T(r, Y_T) = (y_{T,1} + rT^{1/2}, \dots, y_{T,T} + rT^{1/2})'$, $\phi_T(Y_T) = g_T(-T^{-1/2}\bar{y}_T, Y_T)$, and $\tilde{\rho}(x) = \int_0^1 x(s)ds$, we find $(h_T \circ \phi_T)(Y_T) = T^{-1/2}\hat{u}_{T, \lfloor \cdot T \rfloor} \rightsquigarrow J_c^\mu \sim \tilde{\phi}(J_c)$, $\sup_{y \in \mathbb{R}^T} d_{D_{[0,1]}}(T^{-1/2}(y_{\lfloor sT \rfloor} + T^{1/2}r), T^{-1/2}y_{\lfloor sT \rfloor} + r) = 0$, and $x = \tilde{\phi}(x) + \tilde{\rho}(x) = \tilde{g}(\tilde{\rho}(x), \tilde{\phi}(x))$ for all $x \in D_{[0,1]}$, so that the assumptions of Theorem 2 hold. We can therefore conclude that rejecting for large values of $L^\mu(T^{-1/2}\hat{u}_{T, \lfloor \cdot T \rfloor})$ is also the asymptotically point-optimal test among all translation invariant tests with asymptotic null rejection probability of at most α whenever $T^{1/2}y_{T, \lfloor \cdot T \rfloor} \rightsquigarrow J_0(\cdot) + \alpha_y$, $\alpha_y \in \mathbb{R}$. \blacktriangle

For the proof of Theorem 2, note that with $x = \tilde{g}(\tilde{\rho}(x), \tilde{\phi}(x))$ for all $x \in S$ and $\theta \in \Theta$, one can construct the distribution $P(\theta)$ by applying an appropriate *random* transformation \tilde{g} to each x drawn under $\tilde{\phi}P(\theta)$. The assumptions of Theorem 2 are now sufficient to ensure a tight enough link between this construction for $P(\theta)$ and the limit of the small sample analogously transformed \mathcal{M}_j^ϕ models. For each model in \mathcal{M}_j^ϕ , one can thus construct a corresponding model in \mathcal{M}_j by applying an appropriate random transformation $g_T \in \mathcal{G}_T$ for each T .

While asymptotic efficiency statements about invariant tests based on Theorems 1 and 2 require a tight link between the limiting group $\tilde{\mathcal{G}}$ and the small sample groups \mathcal{G}_T , the link does not need to be perfect: Even if the distribution of the limiting maximal invariant $\tilde{\phi}P(\theta)$ does not depend a subset of the parameter θ (so that a nuisance parameter is eliminated by invariance), it is not assumed that the small sample counterpart $\phi_T F_T(m, \theta)$ shares this feature. Also assumption (iv) does not require the small sample and limiting group actions to exactly coincide, which is important for, say, arguing for the asymptotic efficiency

of translation and trend invariant unit root tests with respect to the weak convergence $T^{-1/2}y_{T,[\cdot T]} \rightsquigarrow J_c(\cdot) + \alpha_y + \cdot\beta_y$.

3.3 Uniformity

The discussion so far concerned the pointwise asymptotic properties of tests φ_T , i.e. the rejection probability as $T \rightarrow \infty$ for a fixed model m and parameter value θ . This is standard practice in much of econometric theory, including in the literature on semiparametrically efficient tests.

When one insists on uniform results in the set-up here, it is natural to ask whether the pointwise robustness constraint (8) can be replaced by a uniform constraint over models, that is to demand that for large enough T , the null rejection probability $\int \varphi_T dF_T(m, \theta)$ for $\theta \in \Theta_0$ to be close to α for all data generating processes under consideration. It is not hard to see, however, that with \mathcal{M}_0 the set of all models satisfying (2), such a uniformity cannot hold for non-trivial tests: for any T , the distribution $F_T(m, \theta)$ of Y_T is entirely unrestricted, as the convergence $P_T(m, \theta) = h_T F_T(m, \theta) \rightsquigarrow P(\theta)$ can occur 'later'.

To generate uniform results, one must therefore reduce the set of models \mathcal{M}_0 and impose a lower limit on the speed of convergence. For two probability measures μ_1 and μ_2 on $\mathfrak{B}(S)$, define Δ_{BL} as $\Delta_{BL}(\mu_1, \mu_2) = \sup_{\|f\|_{BL} \leq 1} |\int f d\mu_1 - \int f d\mu_2|$, where $f : S \mapsto \mathbb{R}$ are $\mathfrak{B}(S) \setminus \mathfrak{B}(\mathbb{R})$ measurable and $\|f\|_{BL} = \sup_{x \in S} |f(x)| + \sup_{x \neq y, x, y \in S} \frac{|f(x) - f(y)|}{d_S(x, y)}$. It is known that Δ_{BL} metrizes weak convergence on separable metric spaces (Dudley (2002, p. 395)). Also, let the real sequence δ_T be such that $\delta_T \rightarrow 0$. Now define $\mathcal{M}_0^u(\delta)$ ('u' for uniform) as the set of models m satisfying

$$\sup_{\theta \in \Theta_0} \Delta_{BL}(P_T(m, \theta), P(\theta)) \leq \delta_T,$$

that is $\mathcal{M}_0^u(\delta)$ is the collection of models m for which the distribution $P_T(m, \theta) = h_T F_T(m, \theta)$ of $h_T(Y_T)$ differs by at most δ_T from its limit $P(\theta)$ as measured by Δ_{BL} , uniformly over Θ_0 . It then makes sense to ask whether the rejection probability of a test φ_T converges to the nominal level uniformly over $\theta \in \Theta_0$ and $\mathcal{M}_0^u(\delta)$, that is if

$$\limsup_{T \rightarrow \infty} \sup_{\theta \in \Theta_0, m \in \mathcal{M}_0^u(\delta)} \int \varphi_T dF_T(m, \theta) \leq \alpha. \quad (18)$$

By the continuity of φ_S^* , (18) holds for the test $\varphi_T = \hat{\varphi}_T^*$ in Theorem 1, i.e. for large enough

T , the rejection probability of $\hat{\varphi}_T^*$ is close to α for all models in $\mathcal{M}_0^u(\delta)$, uniformly over Θ_0 .⁴ Similar restrictions and arguments could be made regarding the constraints (10) and (11).

It is not clear, however, whether all tests that satisfy the point-wise robustness (8) also satisfy (18), or vice versa. Theorem 1 therefore does not imply that $\hat{\varphi}_T^*$ also maximizes asymptotic weighted average power in the class of all tests that satisfy (18). A partial result in that regard is provided by the following Theorem for the case of a single null hypothesis and without the constraints (10) and (11).

Theorem 3 *Suppose $\Theta_0 = \{\theta_0\}$, let $L : \Theta \times S \mapsto \mathbb{R}$ be the Radon-Nikodym derivative of $P(\theta)$ with respect to $P(\theta_0)$, define $\bar{L}(x) = \int L(\theta, x)dw(\theta)$, let $\varphi_S^* : S \mapsto [0, 1]$ be the level α test that rejects for large values of \bar{L} , and define $\hat{\varphi}_T^* = \varphi_S^* \circ h_T$. Suppose that for all $\varepsilon > 0$ there exists an open set $D_\varepsilon \in \mathfrak{B}(S)$ with $\int_{D_\varepsilon} dP(\theta_0) > 1 - \varepsilon$ so that the $D_\varepsilon \mapsto \mathbb{R}$ function $x \mapsto \bar{L}(x)$ is Lipschitz, and assume that the models m_0 and m_1 are such that $\Delta_{BL}(P_T(m_0, \theta_0), P(\theta_0))/\delta_T \rightarrow 0$ and $\Delta_{BL}(\int P_T(m_1, \theta)dw(\theta), \int P(\theta)dw(\theta))/\delta_T \rightarrow 0$. Then for any test φ_T that satisfies (18), $\limsup_{T \rightarrow \infty} \text{WAP}(\varphi_T, m_1) \leq \lim_{T \rightarrow \infty} \text{WAP}(\hat{\varphi}_T^*, m_1) = \text{WAP}(\varphi_S^*)$.*

Under a stronger continuity assumption on the limiting problem, Theorem 3 shows that no test can satisfy (18) and have higher asymptotic weighted average power than $\hat{\varphi}_T^*$ for alternative models whose (average) weak convergence is faster than the lower bound δ_T . In other words, $\hat{\varphi}_T^*$ is the efficient test in the class of test satisfying (18) against alternative model m_1 for which $\int P_T(m_1, \theta)dw(\theta) \rightsquigarrow \int P(\theta)dw(\theta)$ converges faster than $P_T(m, \theta_0) \rightsquigarrow P(\theta_0)$ in (18).

The proof of Theorem 3 follows closely the heuristic sketch of the proof Theorem 1 outlined above and exploits the linearity in both the definition of Δ_{BL} and the probability assignments in the tilted model G_T . Dudley (2002, p. 411) shows that the Prohorov metric Δ_P (which also metrizes weak convergence) satisfies $\Delta_P \leq 2\Delta_{BL}^{1/2}$ and $\Delta_{BL} \leq 2\Delta_P$, so that Theorem 3 could equivalently be formulated in terms of Δ_P .

⁴Construct a decreasing sequence of functions ℓ_T that converge to φ_S^* pointwise $P(\theta_0)$ -almost everywhere with $\delta_T \|\ell_T\|_{BL} \rightarrow 0$, as in chapter 7.1 of Pollard (2002). Then $\sup_{\theta \in \Theta_0, m \in \mathcal{M}_0^u(\delta)} \int \hat{\varphi}_T^* dF_T(m, \theta) - \alpha \leq \sup_{\theta \in \Theta_0, m \in \mathcal{M}_0^u(\delta)} \int \ell_T(dP_T(m, \theta) - dP(\theta)) \leq \|\ell_T\|_{BL} \sup_{\theta \in \Theta_0} \Delta_{BL}(P_T(m, \theta), P(\theta)) \rightarrow 0$, where the last inequality uses $\|f_1 \cdot f_2\|_{BL} \leq \|f_1\|_{BL} \cdot \|f_2\|_{BL}$.

4 Applications

4.1 Unit Root Tests with Stationary Covariates

Elliott and Jansson (2003) consider the model

$$\begin{pmatrix} y_{T,t} \\ x_{T,t} \end{pmatrix} = \begin{pmatrix} \alpha_y + \beta_y t \\ \alpha_x + \beta_x t \end{pmatrix} + \begin{pmatrix} u_{T,t} \\ \nu_{T,t}^x \end{pmatrix} \quad (19)$$

where $Y_T = ((y_{T,1}, x'_{T,1})', \dots, (y_{T,T}, x'_{T,T}))' \in \mathbb{R}^{nT}$ is observed, α_y, β_y and $u_{T,t} = \rho_T u_{T,t-1} + \nu_{T,t}^y$ are scalars, $u_{T,0} = O_p(1)$, $\rho_T = 1 - c/T$ for some fixed $c \geq 0$, and $x_{T,t}, \alpha_x, \beta_x$ and $\nu_{T,t}^x$ are $n - 1$ dimensional vectors. The objective is to efficiently exploit the stationary covariates $x_{T,t}$ in the construction of a test of the null hypothesis of a unit root in $y_{T,t}$, $H_0 : c = 0$ against the alternative $H_1 : c > 0$. Consider first the case with $\alpha_y = \alpha_x = \beta_y = \beta_x = 0$ known. The approach of Elliott and Jansson (2003) is to first apply the Neyman-Pearson Lemma to determine, for each T , the point-optimal test against $c = c_1$ when $\nu_{T,t} = (\nu_{T,t}^y, \nu_{T,t}^x)' \sim \text{i.i.d.} \mathcal{N}(0, \Omega)$ for known Ω . In a second step, they construct a feasible test that is (i) asymptotically equivalent to the point-optimal test when $\nu_{T,t} \sim \text{i.i.d.} \mathcal{N}(0, \Omega)$ and (ii) that is robust to a range of autocorrelation structures and error distributions. So by construction, their test can only claim efficiency for the special case of i.i.d. Gaussian disturbances.

In order to apply the results in Sections 2 and 3 of this paper, we consider the typical weak convergence properties of model (19). Standard weak dependence assumptions on $\nu_{T,t}$ imply for some suitable long-run covariance matrix estimator $\hat{\Omega}_T$

$$\hat{\Omega}_T \xrightarrow{p} \Omega \quad \text{and} \quad G_T(\cdot) = \begin{pmatrix} T^{-1/2} u_{T, [\cdot T]} \\ T^{-1/2} \sum_{t=1}^{[\cdot T]} \nu_{T,t}^x \end{pmatrix} \rightsquigarrow G(\cdot) \quad (20)$$

where Ω is positive definite, $G(s) = \int_0^s \text{diag}(e^{-c(s-r)}, 1, \dots, 1) \Omega^{1/2} dW(r)$, and W is a $n \times 1$ standard Wiener process. By Girsanov's Theorem, the Radon-Nikodym derivative of the distribution of G with $c = c_1$ with respect to the distribution of G with $c = 0$, evaluated at $G = (G_y, G_x)'$, is given by

$$\begin{aligned} L(\Omega, G) &= \exp \left[-c_1 \int_0^1 G(s)' S_1 \Omega^{-1} dG(s) - \frac{1}{2} c_1^2 \int_0^1 G(s)' S_1 \Omega^{-1} S_1 G(s) ds \right] \\ &= \exp \left[-\frac{1}{2} c_1 \omega_{yy} (G_y(1))^2 - \sigma_{yy}^2 - c_1 \omega_{yx} \int_0^1 G_y(s) dG_x(s) - \frac{1}{2} c_1^2 \omega_{yy} \int_0^1 G_y(s)^2 ds \right] \end{aligned} \quad (21)$$

where S_1 is the $n \times n$ matrix $S_1 = \text{diag}(1, 0, \dots, 0)$, σ^2 is the (1,1) element of Ω , and the first row of Ω^{-1} is $(\omega_{yy}, \omega_{yx})$. By the Neyman-Pearson Lemma, the point-optimal test in the limiting problem rejects for large values of $L(\Omega, G)$. Since $\int_0^1 G_y(s) dG_x(s)$ is not a continuous mapping, we cannot directly apply Theorem 1. However, typical weak dependence assumptions on $\nu_{T,t}$ also imply (see, for instance, Phillips (1988), Hansen (1990) and de Jong and Davidson (2000)) that

$$\Upsilon_T = T^{-1} \sum_{t=2}^T u_{T,t-1} \nu_{T,t}^x - \hat{\Sigma}_T \rightsquigarrow \Upsilon = \int_0^1 G_y(s) dG_x(s) \quad (22)$$

for a suitably defined $(n-1) \times 1$ vector $\hat{\Sigma}_T \xrightarrow{p} \Sigma$ (which equals $\sum_{s=1}^{\infty} E[\nu_{T,t}^x \nu_{T,t+s}^y]$ when $\nu_{T,t}$ is covariance stationary) jointly with (20). Clearly, the Radon-Nikodym derivative of the measure of (G, Υ) for $c = c_1$ with respect to the measure of (G, Υ) with $c = 0$, evaluated at G , is also given by $L(\Omega, G)$ in (21), and one can write $L(\Omega, G) = L^{\Upsilon}(\Omega, G, \Upsilon)$ for a continuous function L^{Υ} . The discussion of Section 3.1 thus applies, and Theorem 1 shows that rejecting for large values of $L^{\Upsilon}(\hat{\Omega}_T, G_T, \Upsilon_T)$ is the point-optimal unit root test for the alternative $c = c_1$ among all tests that have correct asymptotic null rejection probabilities whenever (20) and (22) hold.

Since the model with $\nu_{T,t} \sim \text{i.i.d.} \mathcal{N}(0, \Omega)$ satisfies (20) and (22), the test derived by Elliott and Jansson (2003) is by construction—as explained in Comment 1 of Section 2.5— asymptotically equivalent to a test that rejects for large values of $L^{\Upsilon}(\hat{\Omega}_T, G_T, \Upsilon_T)$. The derivation here, which starts with the Radon-Nikodym derivative directly, is arguably a more straightforward way of determining a test in this equivalence class. But the more important insight concerns the optimality properties of this test: While Elliott and Jansson (2003) could only claim optimality for the model with i.i.d. Gaussian disturbances, Theorem 1 shows that the test to be efficient against *all* alternatives satisfying (20) and (22) with $c = c_1$ if one imposes size control for all models satisfying (20) and (22) with $c = 0$. In other words, under this robustness constraint, no test exists with higher asymptotic power for any disturbance distribution or autocorrelation structure satisfying (20) and (22) with $c = c_1$.

When the deterministic terms are not fully known, i.e. the parameters α_y , α_x , β_y , and/or β_x are not known, it is natural to impose an appropriate invariance requirement. Specifically, considering the case where α_y and α_x are unconstrained and $\beta_y = \beta_x = 0$, one

might impose invariance to the transformations

$$\{(y_{T,t}, x'_{T,t})'\}_{t=1}^T \rightarrow \{(y_{T,t} + a_y, x'_{T,t} + a'_x)\}_{t=1}^T \quad a_y \in \mathbb{R}, a_x \in \mathbb{R}^{n-1}. \quad (23)$$

A maximal invariant of this group of transformations is given by the demeaned data $\{(\hat{y}_{T,t}, \hat{x}'_{T,t})'\}_{t=1}^T$, where $\hat{y}_{T,t} = y_{T,t} - \bar{y}_T$, $\hat{x}_{T,t} = x_{T,t} - \bar{x}_T$, $\bar{y}_T = T^{-1} \sum_{t=1}^T y_{T,t}$ and $\bar{x}_T = T^{-1} \sum_{t=1}^T x_{T,t}$. Elliott and Jansson (2003) derive the limiting behavior of the likelihood ratio statistics of this maximal invariant when $\nu_{T,t} \sim \text{i.i.d.} \mathcal{N}(0, \Omega)$, and thus obtain the asymptotically point-optimal invariant unit root test under that assumption. Considering again the weak convergence properties of a typical model, we obtain

$$\hat{\Omega}_T \xrightarrow{p} \Omega \quad \text{and} \quad \begin{pmatrix} T^{-1/2} \hat{y}_{T, \lfloor \cdot T \rfloor} \\ T^{-1/2} \sum_{t=1}^{\lfloor \cdot T \rfloor} \hat{x}_{T,t} \\ T^{-1/2} \sum_{t=2}^T \hat{y}_{T,t-1} \hat{x}_{T,t} - \hat{\Sigma}_T \end{pmatrix} \rightsquigarrow \begin{pmatrix} \hat{G}_y(\cdot) \\ \hat{G}_x(\cdot) \\ \int_0^1 \hat{G}_y(s) d\hat{G}_x(s) \end{pmatrix} \quad (24)$$

where $\hat{G}_y(s) = G_y(s) - \int_0^1 G_y(l) dl$ and $\hat{G}_x(s) = G_x(s) - sG_x(1)$. For brevity, we omit an explicit expression for the Radon-Nikodym derivative $L^{\hat{G}}$ of the measure of (\hat{G}_y, \hat{G}_x) with $c = c_1$ with respect to the measure of (\hat{G}_y, \hat{G}_x) when $c = 0$. By the Neyman-Pearson Lemma and Theorem 1, rejecting for large values of $L^{\hat{G}}$, evaluated at sample analogues, is the asymptotically point-optimal robust test for models defined via (24). Furthermore, in the notation of Section 3.2, with $h_T(Y_T) = (T^{-1/2} y_{T, \lfloor \cdot T \rfloor}, T^{-1/2} \sum_{t=1}^{\lfloor \cdot T \rfloor} x_{T,t}, T^{-1} \sum_{t=2}^T y_{T,t-1} x_{T,t} - \hat{\Sigma}_T) \in D_{[0,1]} \times D_{[0,1]}^{n-1} \times \mathbb{R}^{n-1}$, $r = (r_y, r'_x)' \in \mathbb{R} \times \mathbb{R}^{n-1}$, $g_T(r, \{(y_{T,t}, x'_{T,t})'\}_{t=1}^T) = \{(y_{T,t} + T^{1/2} r_y, x'_{T,t} + T^{-1/2} r'_x)\}'_{t=1}^T$, $\tilde{g}(r, (y, x, z)) = (y(\cdot) + r_y, x(\cdot) + r_x \cdot, z + r_y(x(1) - x(0)) + r_x \int_0^1 y(s) ds + r_y r_x)$, $\phi_T(\{(y_{T,t}, x'_{T,t})'\}_{t=1}^T) = \{(\hat{y}_{T,t}, \hat{x}'_{T,t})'\}_{t=1}^T$ and $\tilde{\phi}(y, x, z) = (y(\cdot) - \int_0^1 y(s) ds, x(\cdot) - \cdot x(1), z - (x(1) - x(0)) \int_0^1 y(s) ds)$, we find that Theorem 2 is applicable, and that rejecting for large values of $L^{\hat{G}}$, evaluated at sample analogues, is also the asymptotically point-optimal invariant unit root test among all tests with correct asymptotic null rejection probability for all models satisfying (20) and (22) with $c = 0$.

4.2 Linear Regression with Weak Instruments

As Andrews, Moreira, and Stock (2006) (abbreviated AMS in the following), consider the problem of inference about the coefficient of a scalar endogenous variable in the presence of weak instruments. The reduced form equations are given by (cf. equation (2.4) of AMS)

$$\begin{aligned} y_{1,t} &= z'_t \pi \beta + x'_t \zeta_1 + v_{t,1} \\ y_{2,t} &= z'_t \pi + x'_t \zeta_2 + v_{t,2} \end{aligned} \quad (25)$$

for $t = 1, \dots, T$, where $y_{1,t}$ and $y_{2,t}$ are scalars, z_t is $k \times 1$ and x_t is $p \times 1$, and z_t are the residuals of a linear regression of the original instruments \tilde{z}_t on x_t . AMS initially consider small sample efficient tests of

$$H_0 : \beta = \beta_0 \quad (26)$$

for nonstochastic (x_t, z_t) and $v_t = (v_{1,t}, v_{2,t})' \sim \text{i.i.d.} \mathcal{N}(0, \Omega)$ with Ω known. A sufficiency argument shows that tests may be restricted to functions of the $2k \times 1$ multivariate normal statistic

$$\sum_{t=1}^T \begin{pmatrix} z_t y_{1,t} \\ z_t y_{2,t} \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} S_z \pi \beta \\ S_z \pi \end{pmatrix}, \Omega \otimes S_z \right) \quad (27)$$

where $S_z = \sum_{t=1}^T z_t z_t'$, and AMS derive weighted average power maximizing similar tests that are invariant to the group of transformations

$$\{z_t\}_{t=1}^T \rightarrow \{O z_t\}_{t=1}^T \text{ for any orthogonal matrix } O. \quad (28)$$

For their asymptotic analysis, AMS employ Staiger and Stock (1997) weak instrument asymptotics, where $\pi = T^{-1/2} C$ for some fixed matrix C . AMS then exploit their small sample efficiency results to construct tests that (i) maximize weighted average asymptotic power among all asymptotically invariant and asymptotically similar test when $v_t \sim \text{i.i.d.} \mathcal{N}(0, \Omega)$ independent of $\{(x'_t, z'_t)\}_{t=1}^T$; and (ii) that yield correct null rejection probability under much broader conditions (the working paper of AMS contains the details of the construction of heteroskedasticity, and of heteroskedasticity and autocorrelation robust tests).

To apply the results from Sections 2 and 3 to this problem, consider the set of weak convergences in the double array version of model (25)

$$\begin{aligned} \hat{D}_Z &= T^{-1} \sum_{t=1}^T z_{T,t} z'_{T,t} \xrightarrow{p} D_z, \hat{\Sigma} \xrightarrow{p} \Sigma, \\ X_T &= T^{-1/2} \sum_{t=1}^T \begin{pmatrix} z_{T,t} y_{1,T,t} \\ z_{T,t} y_{2,T,t} \end{pmatrix} \rightsquigarrow X \sim \mathcal{N} \left(\begin{pmatrix} D_z C \beta \\ D_z C \end{pmatrix}, \Sigma \right) \end{aligned} \quad (29)$$

where $\hat{\Sigma}$ is some standard estimator of the long run variance of $\text{vec}(z_t v'_t)$, and D_z and Σ have full rank. The limiting problem in the sense of Section 2.2 above is thus the test of (26) based on observing the random variable X distributed as in (29), with D_z and Σ known.

For $k = 1$, i.e. in the just-identified case, this problem is exactly equivalent to the small sample problem (27) considered by AMS. In fact, already Moreira (2001) has shown that for $k = 1$, the Anderson and Rubin (1949) statistic $AR_0 = (b'_0 X)^2 / b'_0 \Sigma b_0$ with $b_0 = (1, -\beta_0)'$

yields the uniformly most powerful unbiased test φ_S^* . The discussion of Section 3.1 and Theorem 1 thus imply that for $k = 1$, rejecting for large values of $(b_0'X_T)^2/b_0'\hat{\Sigma}b_0$ maximizes asymptotic power uniformly in all models that satisfy (29) with $\beta \neq \beta_0$ among all tests that are asymptotically unbiased with at most nominal asymptotic null rejection probability for all models that satisfy (29) with $\beta = \beta_0$.

The robustness constraint here—correct asymptotic null rejection probability for all models that satisfy (29) with $\beta = \beta_0$ —is quite stringent, since there are many ways $T^{-1/2} \sum_{t=1}^T z_{T,t}y_{T,t}$ can converge to a normal vector (for instance, the convergence is compatible with $z_{T,t}y_{T,t} = 0$ for all $t < T/2$). To the extent that one would be prepared to rule out such models a priori, this decreases the appeal of the efficiency result. But as discussed in comment 6 in Section 2.5, one can impose additional weak convergences without necessarily affecting φ_S^* : Supplementing (29) by the FCLT type convergence

$$T^{-1/2} \sum_{t=1}^{\lfloor T \rfloor} \begin{pmatrix} z_{T,t}y_{1,T,t} \\ z_{T,t}y_{2,T,t} \end{pmatrix} \rightsquigarrow G(\cdot) \text{ with } G(s) = s \begin{pmatrix} D_z C \beta \\ D_z C \end{pmatrix} + \Sigma^{1/2}W(s),$$

with $W(s)$ is a $2k \times 1$ standard Wiener process, for instance, rules out such pathological cases, and yet still yields φ_S^* to be the efficient test, since $G(1)$ is sufficient for the unknown parameters C and β .

For $k > 1$, Theorems 1 and 2 can again be invoked to yield analogous asymptotic efficiency statements for the statistics developed in AMS under the assumption that Σ is of the Kronecker form $\Sigma = \Omega \otimes D_z$, as in (27). But this form naturally arises only in the context of a serially uncorrelated homoskedastic model, so the resulting efficiency statements are of limited appeal. The approach here thus points to a general solution of the limiting problem without the constraint $\Sigma = \Omega \otimes D_z$ as an interesting missing piece in the literature on efficient inference in linear regressions with weak instruments.

4.3 GMM Parameter Stability Tests

Following Sowell (1996), suppose we are interested in testing the null hypothesis that a parameter $\beta \in \mathbb{R}^k$ in a GMM framework is constant through time. Parametrizing $\beta_{T,t} = \beta_0 + T^{-1/2}\theta(t/T)$, where $\theta \in D_{[0,1]}^k$ and $\theta(0)$ is normalized to zero $\theta(0) = 0$, this is equivalent to the hypothesis test

$$H_0 : \theta = 0 \quad \text{against} \quad H_1 : \theta \neq 0. \tag{30}$$

Denote by $g_{T,t}(\beta) \in \mathbb{R}^p$ with $p \geq k$ the sample moment condition for $y_{T,t}$ evaluated at β , so that under the usual assumptions, the moment condition evaluated at the true parameter value satisfies a central limit theorem, that is $T^{-1/2} \sum_{t=1}^T g_{T,t}(\beta_{T,t}) \rightsquigarrow \mathcal{N}(0, V)$ for some positive definite $p \times p$ matrix V . Furthermore, with $\hat{\beta}_T$ the usual full sample GMM estimator of β with optimal weighting matrix converging to V^{-1} , we obtain under typical assumptions that for some suitable estimators \hat{H}_T and \hat{V}_T (cf. Theorem 1 of Sowell (1996))

$$G_T(\cdot) = T^{-1/2} \sum_{t=1}^{\lfloor T \rfloor} g_{T,t}(\hat{\beta}_T) \rightsquigarrow G(\cdot) \quad \text{and} \quad \hat{H}_T \xrightarrow{p} H, \hat{V}_T \xrightarrow{p} V \quad (31)$$

where the convergence to G is on $D_{[0,1]}^p$, $G(s) = V^{1/2}W(s) - sH(H'V^{-1}H)^{-1}H'V^{-1/2}W(1) + H \left(\int_0^s \theta(l)dl - s \int_0^1 \theta(l)dl \right)$ with W a $p \times 1$ standard Wiener process and H some $p \times k$ matrix full column rank matrix (which is the probability limit of the average of the partial derivatives of $g_{T,t}$). Andrews (1993), Sowell (1996) and Li and Müller (2007) discuss primitive conditions for these convergences. Sowell (1996) goes on to derive weighted average power maximizing tests of (30) as a function of G (that is, he computes φ_S^* in the notation of Theorem 1), and he denotes the resulting test evaluated at $G_T(\cdot)$, \hat{H}_T and \hat{V}_T (that is, $\hat{\varphi}_T^*$ in the notation of Theorem 1), an "optimal" test for structural change.

Without further restrictions, however, such tests cannot claim to be efficient: As a simple example, consider the scalar model with $y_{T,t} = \beta + T^{-1/2}\theta(t/T) + \varepsilon_t$, where ε_t is i.i.d. with $P(\varepsilon_t = -1) = P(\varepsilon_t = 1) = 1/2$. This model is a standard time varying parameter GMM model with $g_{T,t}(\beta_0) = y_{T,t} - \beta_0 = T^{-1/2}\theta(t/T) + \varepsilon_t$ satisfying (31), yet in this model, the test φ_T^{**} that rejects whenever any one of $\{y_{T,t} - y_{T,t-1}\}_{t=1}^T$ is not $-2, 0$ or 2 has level zero for any $T \geq 2$ and has asymptotic power equal to one against any local alternative.

Theorem 1 provides a sense in which the tests derived by Sowell (1996) are asymptotically optimal: they maximize asymptotic weighted average power among all tests that have correct asymptotic null rejection probability whenever (31) holds with $\theta = 0$. More powerful tests that exploit specificities of the error distribution, such as φ_T^{**} , necessarily do not lead to correct asymptotic null rejection probability for all stable models satisfying (31).

5 Conclusion

This paper analyzes an alternative notion of asymptotic efficiency of tests. The starting point is the idea that models of interest can be described in terms of a weak convergence

property, so that under the null and alternative hypothesis, an observed random element converges weakly to a limiting random element. It is shown that if one restricts attention to tests that have nominal asymptotic rejection probability for all models that satisfy the weak convergence under the null hypothesis, then efficient tests in the original problem are simply given by efficient tests in the limiting problem (that is, with the limiting random element observed), evaluated at sample analogues. These efficient tests generically coincide with robustified versions of efficient tests that are derived as the limit of small sample efficient tests in canonical parametric versions of the model. The results of this paper thus provide an alternative and broader sense of asymptotic efficiency for many previously derived tests in econometrics.

It is a strong restriction to force tests to have correct asymptotic null rejection probability for all models that satisfy the weak convergence, because there are typically many such models. Relative to the traditional semiparametrically efficient testing literature, the issue of asymptotic efficiency is thus approached here with a stronger focus on robustness of inference. This allows for a general and easily applied efficiency claim under minimal additional regularity conditions. Given the focus on the robustness, the results derived here are most naturally applied in time series econometrics, where there often is substantial uncertainty over the joint distribution of the data, and where traditional semiparametric considerations face serious challenges. The two approaches thus naturally complement each other.

6 Appendix

Proof of Theorem 1:

(i) We prove the second claim, the first is proved analogously.

Since φ_S^* is $P(\theta_0)$ -almost everywhere continuous, it is also $P(\theta)$ -almost everywhere continuous for any $\theta \in \Theta_1$, and by the Continuous Mapping Theorem, $P_T(m, \theta) \rightsquigarrow P(\theta)$ implies $\int \varphi_S^* dP_T(m, \theta) \rightarrow \int \varphi_S^* dP(\theta)$ for all $\theta \in \Theta_1$. Since $0 \leq \varphi_S^* \leq 1$, the result follows by dominated convergence.

(ii) Heuristically, the proof follows the basic logic of constructing a tilted probability measure from a given alternative model described in the main text, but it addresses the difficulty that (a) L^i is not necessarily bounded; and (b) the testing problem is more general than a simple null against a simple alternative, and tests are further constrained by similarity and unbiasedness conditions.

The proof consists of three steps. Consider first complication (a). Note that in the limit where $X \sim P(\theta)$, any change of measure from $P(\theta_1)$ to $P(\theta_0)$ using $L = dP(\theta_0)/dP(\theta_1)$, $\int_A dP(\theta_0) = \int_A \frac{1}{L} dP(\theta_1) = \int_A \frac{1}{L} L dP(\theta_0)$ for all $A \in \mathfrak{B}(S)$, is well defined, even if L^i is unbounded. The idea of the proof is thus to perform the tilting 'in the limit' (Step 2), and ensure that it has the desired weak convergence implications for X_T by constructing a probability space in which $X_T = h(Y_T) \rightarrow X$ almost surely (Step 1). Such almost sure representations of weakly converging random elements are known to exist in a complete and separable space. Because X_T is a function of Y_T , and the set of models \mathcal{M}_0 is defined relative to the distribution of Y_T , we must use a construction that preserves the information about the distribution of Y_T . The tilting of the limiting measure in Step 2 then also implies the appropriate change of measure for $Y_T = h_T(X_T)$, as demonstrated in Step 3.

Due to complication (b), we cannot directly apply the Neyman-Pearson Lemma to learn about the efficiency of a candidate test φ_T for distinguishing between the original alternative model and the tilted model. Rather, we compare φ_T with φ_S^* directly 'in the limit' in Step 2 by making its (best) asymptotic performance under the alternative part of the almost surely converging construction in Step 1. This is possible, because (randomized) tests take on values in the unit interval, so that any sequence of φ_T has a subsequence that converges weakly by Prohorov's Theorem. Because the tilting in the limit is a function of X only, X is a sufficient statistic, so that $\varphi_S^*(X)$ is an overall best test for learning about the tilting—see Step 2. Finally, as φ_T is known to satisfy (8), (10) and (11), its weak limit is a level α test

and satisfies (6) and (7) under the tilted probability measures, so that it cannot be better than φ_S^* —see Step 3.

Step 1: Pick any $m \in \mathcal{M}_1$. Define $\bar{F}_T = \int F_T(m, \theta)dw(\theta)$, $\bar{P}_T = \int P_T(m, \theta)dw(\theta)$ and $\bar{P} = \int P(\theta)dw(\theta)$. For any bounded and continuous function $\vartheta : S \mapsto \mathbb{R}$, $\int \vartheta dP_T(m, \theta) \rightarrow \int \vartheta dP(\theta)$ for all $\theta \in \Theta_1$ by the continuous mapping theorem, so that by dominated convergence, also $\int \vartheta d\bar{P}_T \rightarrow \int \vartheta d\bar{P}$. Thus, $h_T \bar{F}_T = \bar{P}_T \rightsquigarrow \bar{P}$. Let $D_{[0,1]}^n$ be the space of n -valued cadlag functions on the unit interval, equipped with the Billingsley (1968) metric, and define the mapping $\chi_T : \mathbb{R}^{nT} \mapsto D_{[0,1]}^n$ as $\{y_t\}_{t=1}^T \mapsto T^{-1}\Phi_Z(y_{[\cdot, T]})$, where Φ_Z is the c.d.f. of a standard normal applied element by element. Note that χ_T is injective, and denote by χ_T^{-1} a $D_{[0,1]}^n \mapsto \mathbb{R}^{nT}$ function such that $\chi_T^{-1}(\chi_T(y)) = y$ for all $y \in \mathbb{R}^{nT}$. Since $\sup_{s \in [0,1]} \|\chi_T(s)\| \leq 1/T \rightarrow 0$, the probability measures $(h_T, \chi_T)\bar{F}_T$ on the complete and separable space $S \times D_{[0,1]}^n$ converge weakly to the product measure $\bar{P} \times \delta_0$, where δ_0 puts all mass at the zero function in $D_{[0,1]}^n$. Let $T_1 \rightarrow \infty$ be any subsequence of T such that $\lim_{T_1 \rightarrow \infty} \int \varphi_{T_1} d\bar{F}_{T_1} = \limsup_{T \rightarrow \infty} \int \varphi_T d\bar{F}_T$. Since the probability measures $(h_{T_1}, \chi_{T_1}, \varphi_{T_1})\bar{F}_{T_1}$ on the complete and separable space $S \times D_{[0,1]}^n \times [0, 1]$ are uniformly tight, by Prohorov's Theorem (see, for instance, Theorem 36 on p. 185 of Pollard (2002)), there exists a subsequence T_2 of T_1 such that $(h_{T_2}, \chi_{T_2}, \varphi_{T_2})\bar{F}_{T_2} \rightsquigarrow \bar{\nu}$ as $T_2 \rightarrow \infty$, where $(\pi_X, \pi_Y)\bar{\nu} = \bar{P} \times \delta_0$, and π_X , π_Y and π_φ are the projections of $S \times D_{[0,1]}^n \times [0, 1]$ on S , $D_{[0,1]}^n$ and $[0, 1]$, respectively. For notational convenience, write T for T_2 in the following. By Theorem 11.7.2 of Dudley (2002), there exists a probability space $(\Omega^*, \mathcal{F}^*, P^*)$ and functions $\eta_T : \Omega^* \mapsto S \times D_{[0,1]}^n \times [0, 1]$ such that $\eta_T P^* = (h_T, \chi_T, \varphi_T)\bar{F}_T$, $\eta_0 P^* = \bar{\nu}$ and $\eta_T(\omega^*) \rightarrow \eta_0(\omega^*)$ for P^* -almost all $\omega^* \in \Omega^*$, and by Theorem 11.7.3 of Dudley (2002), we may assume Ω^* to be complete and separable. In this construction, note that for P^* -almost all ω^* , $h_T \circ \chi_T^{-1} \circ \pi_Y \circ \eta_T(\omega^*) = \pi_X \circ \eta_T(\omega^*)$ and $\varphi_T \circ \chi_T^{-1} \circ \pi_Y \circ \eta_T(\omega^*) = \pi_\varphi \circ \eta_T(\omega^*)$. Furthermore, $(\pi_X \circ \eta_T)P^* = \bar{P}_T$, $(\chi_T^{-1} \circ \pi_Y \circ \eta_T)P^* = \bar{F}_T$ and $\int \varphi_T d\bar{F}_T = \int (\pi_\varphi \circ \eta_T) dP^* \rightarrow \int (\pi_\varphi \circ \eta_0) dP^*$, where the convergence follows from the dominated convergence theorem, since $\pi_\varphi \circ \eta_T(\omega^*) \in [0, 1]$ for all $\omega^* \in \Omega^*$. We need to show that $\int (\pi_\varphi \circ \eta_0) dP^* \leq \int \varphi_S^* d\bar{P}$.

Step 2: Note that $P(\theta)$ is absolutely continuous with respect to \bar{P} for any $\theta \in \Theta$, and denote by $L(\theta)$ the Radon-Nikodym derivative of $P(\theta)$ with respect to \bar{P} , so that for all $A \in \mathfrak{B}(S)$, $\int_A dP(\theta) = \int_A L(\theta) d\bar{P}$ (existence of L is ensured by the Radon-Nikodym Theorem; see, for instance, page 56 of Pollard (2002)). Define $Q^*(\theta)$ to be the probability measure on \mathcal{F}^* , indexed by $\theta \in \Theta$, as $\int_A dQ^*(\theta) = \int_A (L(\theta) \circ \pi_X \circ \eta_0) dP^*$ for all $A \in \mathcal{F}^*$. By construction,

$(\pi_X \circ \eta_0)Q^*(\theta) = P(\theta)$, since for all $A \in \mathfrak{B}(S)$, $\int_A (\pi_X \circ \eta_0)dQ^*(\theta) = \int_A L(\theta)d\bar{P} = \int_A dP(\theta)$.

Consider the hypothesis test

$$H_0^* : \omega^* \sim Q^*(\theta), \theta \in \Theta_0 \quad \text{against} \quad H_1^* : \omega^* \sim Q^*(\theta), \theta \in \Theta_1. \quad (32)$$

Because the Radon-Nikodym derivative between $Q^*(\theta_1)$ and $Q^*(\theta_2)$ is given by $(L(\theta_1)/L(\theta_2)) \circ \pi_X \circ \eta_0$, the statistic $\pi_X \circ \eta_0 : \Omega^* \mapsto S$ is sufficient for θ by the Factorization Theorem (see, for instance, Theorem 2.21 of Schervish (1995)). Thus, for any test $\varphi_\Omega : \Omega^* \mapsto [0, 1]$, one can define a corresponding test $\varphi_S : S \mapsto [0, 1]$ via $\varphi_S(x) = E[\varphi_\Omega | (\pi_X \circ \eta_0)(\omega^*) = x]$, which satisfies $\int \varphi_\Omega dQ^*(\theta) = \int (\varphi_S \circ \pi_X \circ \eta_0) dQ^*(\theta) = \int \varphi_S dP(\theta)$ for all $\theta \in \Theta$ (cf. Theorem 3.18 of Schervish (1995)), and also $\int (f_S \circ \pi_X \circ \eta_0)(\varphi_\Omega - \alpha) dQ^*(\theta) = \int f_S(\varphi_S - \alpha) dP(\theta)$ for any $f_S \in \mathcal{F}_S$. Since the level α test $\varphi_S^* : S \mapsto [0, 1]$ of $H_0 : X \sim P(\theta)$, $\theta \in \Theta_0$ against $H_1 : X \sim P(\theta)$, $\theta \in \Theta_1$ maximizes weighted average power subject to (6) and (7), the level α test $\varphi_S^* \circ \pi_X \circ \eta_0 : \Omega^* \mapsto [0, 1]$ of (32) maximizes weighted average power $\int \int \varphi_\Omega dQ^*(\theta) dw(\theta)$ among all level α tests $\varphi_\Omega : \Omega^* \mapsto [0, 1]$ of (32) that satisfy

$$\int \varphi_\Omega dQ^*(\theta) \geq \pi_0(\theta) \quad \text{for all } \theta \in \Theta_1 \quad (33)$$

$$\int (f_S \circ \pi_X \circ \eta_0)(\varphi_\Omega - \alpha) dQ^*(\theta) = 0 \quad \text{for all } \theta \in \bar{\Theta}_0 \quad \text{and } f_S \in \mathcal{F}_S, \quad (34)$$

and it achieves the same weighted average power $\int \int (\varphi_S^* \circ \pi_X \circ \eta_0) dQ^*(\theta) dw(\theta) = \int \int \varphi_S^* dP(\theta) dw(\theta)$.

Step 3: Now define the sequence of measures $G_T(\theta)$ on $\mathfrak{B}(\mathbb{R}^{nT})$, indexed by $\theta \in \Theta$, via $G_T(\theta) = (\chi_T^{-1} \circ \pi_Y \circ \eta_T)Q^*(\theta)$, which induce the measures $h_T G_T(\theta) = (h_T \circ \chi_T^{-1} \circ \pi_Y \circ \eta_T)Q^*(\theta) = (\pi_X \circ \eta_T)Q^*(\theta)$ on $\mathfrak{B}(S)$. By construction of η_T and absolute continuity of $Q^*(\theta)$ with respect to P^* , we have $\eta_T(\omega^*) \rightarrow \eta_0(\omega^*)$ for $Q^*(\theta)$ -almost all ω^* , and since almost sure convergence implies weak convergence, $h_T G_T(\theta) \rightsquigarrow (\pi_X \circ \eta_0)Q^*(\theta) = P(\theta)$ pointwise in $\theta \in \Theta$. Thus, as φ_T satisfies (8), (10) and (11), we have $\limsup_{T \rightarrow \infty} \int \varphi_T dG_T(\theta) \leq \alpha$ for all $\theta \in \Theta_0$, $\liminf_{T \rightarrow \infty} \int \varphi_T dG_T(\theta) \geq \pi_0(\theta)$ for all $\theta \in \Theta_1$ and $\lim_{T \rightarrow \infty} \int (\varphi_T - \alpha)(f_S \circ h_T) dG_T(\theta) = 0$ for all $\theta \in \bar{\Theta}_0$ and $f_S \in \mathcal{F}_S$. By the dominated convergence theorem and the construction of G_T , $\int \varphi_T dG_T(\theta) = \int (\varphi_T \circ \chi_T^{-1} \circ \pi_Y \circ \eta_T) dQ^*(\theta) = \int (\pi_\varphi \circ \eta_T) dQ^*(\theta) \rightarrow \int (\pi_\varphi \circ \eta_0) dQ^*(\theta)$ for all $\theta \in \Theta$ and $\int (\varphi_T - \alpha)(f_S \circ h_T) dG_T(\theta) = \int ((\pi_\varphi \circ \eta_T) - \alpha)(f_S \circ \pi_X \circ \eta_T) dQ^*(\theta) \rightarrow \int ((\pi_\varphi \circ \eta_0) - \alpha)(f_S \circ \pi_X \circ \eta_0) dQ^*(\theta)$ for all $\theta \in \bar{\Theta}_0$ and $f_S \in \mathcal{F}_S$. We can conclude that the test $\pi_\varphi \circ \eta_0 : \Omega^* \mapsto [0, 1]$ of (32) is of level α and satisfies (33) and (34). Its weighted average power $\int \int (\pi_\varphi \circ \eta_0) dQ^*(\theta) dw(\theta)$ is therefore smaller or equal than the weighted average power

$\int \int \varphi_S^* dP(\theta) dw(\theta)$ of the test $\varphi_S^* \circ \pi_X \circ \eta_0 : \Omega^* \mapsto [0, 1]$. The result now follows from noting that $\int \int (\pi_\varphi \circ \eta_0) dQ^*(\theta) dw(\theta) = \int \int (L(\theta) \circ \pi_X \circ \eta_0) (\pi_\varphi \circ \eta_0) dP^* dw(\theta) = \int (\pi_\varphi \circ \eta_0) dP^*$, since $\int L(\theta) dw(\theta) = 1$ \bar{P} -almost surely, and the change of the order of integration is allowed by Fubini's Theorem.

Proof of Theorem 2:

For notational convenience, write $F_{0,T} = F_T(m, \theta)$, $P_T = h_T F_{0,T}$, $P_{0,T}^\phi = (h_T \circ \phi_T) F_{0,T}$, $P_0 = P(\theta)$ and $P_0^\phi = \tilde{\phi} P_0$. Proceed in analogy to the proof of Theorem 1 part (ii) to argue for the existence of a probability space $(\Omega^*, \mathcal{F}^*, P^*)$ with complete and separable Ω^* and functions $\eta_T : \Omega^* \mapsto S \times D_{[0,1]}^n$ such that $\eta_T P^* = (h_T \circ \phi_T, \chi_T) F_{0,T}$, $\eta_0 P^* = P_0^\phi \times \delta_0$, where the probability measure δ_0 puts all mass on the zero function $D_{[0,1]}^n \mapsto \mathbb{R}$, and $\eta_T(\omega^*) \rightarrow \eta_0(\omega^*)$ for P^* -almost all $\omega^* \in \Omega^*$. In particular, $(\pi_X \circ \eta_T) P^* = P_{0,T}^\phi$ and $(\chi_T^{-1} \circ \pi_Y \circ \eta_T) P^* = F_{0,T}$, where π_X and π_Y are the usual projections of $S \times D_{[0,1]}^n$ on S and $D_{[0,1]}^n$, respectively. Also, for almost all ω^* , $h_T \circ \phi_T \circ \chi_T^{-1} \circ \pi_Y \circ \eta_T(\omega^*) = \pi_X \circ \eta_T(\omega^*)$.

Let ν be the probability measure on $\mathfrak{B}(R \times S)$ induced by $(\tilde{\rho}, \tilde{\phi}) : S \mapsto R \times S$ under P_0 , i.e. $\nu = (\tilde{\rho}, \tilde{\phi}) P_0$. Since $x = \tilde{g}(\tilde{\rho}(x), \tilde{\phi}(x))$ for all $x \in S$, $P_0 = \tilde{g}\nu$. By Proposition 10.2.8 of Dudley (2002) there exists a probability kernel ν_x from $(S, \mathfrak{B}(S))$ to $(R, \mathfrak{B}(R))$ such that for each $A \in \mathfrak{B}(S)$ and $B \in \mathfrak{B}(R)$, $\nu(A \times B) = \int_A \nu_x(B) dP_0^\phi(x)$. Note that the $\Omega^* \times \mathfrak{B}(S) \mapsto [0, 1]$ mapping defined via $(\omega^*, B) \mapsto \nu_{\pi_X \circ \eta_0(\omega^*)}(B)$ is a probability kernel from $(\Omega^*, \mathcal{F}^*)$ to $(R, \mathfrak{B}(R))$. We can thus construct the probability measure μ^* on $(\Omega^* \times R, (\mathcal{F}^* \otimes \mathfrak{B}(R)))$ via $\mu^*(C \times B) = \int_C \nu_{\pi_X \circ \eta_0(\omega^*)}(B) dP^*(\omega^*)$, and by construction, the mapping $(\omega^*, r) \mapsto \tilde{g}(r, \pi_X \circ \eta_0(\omega^*))$ induces the measure P_0 under μ^* .

Now from assumptions (iv), for μ^* -almost all (ω^*, r) , $h_T \circ g_T(r, \phi_T \circ \chi_T^{-1} \circ \pi_Y \circ \eta_T(\omega^*)) = \tilde{g}(r, h_T \circ \phi_T \circ \chi_T^{-1} \circ \pi_Y \circ \eta_T(\omega^*)) + o(1) = \tilde{g}(r, \pi_X \circ \eta_T(\omega^*)) + o(1) \rightarrow \tilde{g}(r, \pi_X \circ \eta_0(\omega^*))$ where the convergence follows from the continuity of \tilde{g} . But almost sure convergence implies weak convergence, so that the measures G_T on $\mathfrak{B}(\mathbb{R}^{nT})$ induced by the mapping $(\omega^*, r) \mapsto g_T(r, \phi_T \circ \chi_T^{-1} \circ \pi_Y \circ \eta_T(\omega^*))$ under μ^* satisfy $h_T G_T \rightsquigarrow P_0$. Finally, from $\phi_T \circ g_T(r, \phi_T(y)) = \phi_T(y)$ for all $r \in R$ and $y \in \mathbb{R}^{nT}$, it follows that $\phi_T G_T = (\phi_T \circ \chi_T^{-1} \circ \pi_Y \circ \eta_T) P^* = \phi_T F_{0,T}$.

Proof of Theorem 3:

The equality follows as in the proof of Theorem 1 part (i).

For the inequality, write $P_0 = P(\theta_0)$, $\bar{P} = \bar{L} P_0$, $F_{0,T} = F_T(m_0, \theta_0)$, $P_{0,T} = h_T F_{0,T}$, $\bar{F}_T = \int F_T(m_1, \theta) dw(\theta)$ and $\bar{P}_T = h_T \bar{F}_T$, so that $P_{0,T} \rightsquigarrow P_0$ and $\bar{P}_T \rightsquigarrow \bar{P}$ by assumption.

Pick $1/2 > \epsilon > 0$ such that $P_0(\bar{L} = \epsilon) = 0$ and define $B_\epsilon = \{x \in S : \bar{L} > \epsilon\}$. Note that $\int_{S \setminus B_\epsilon} d\bar{P} = \int_{S \setminus B_\epsilon} \bar{L} dP_0 \leq \epsilon \int_{S \setminus B_\epsilon} dP_0 \leq \epsilon$, so that $\int_{B_\epsilon} d\bar{P} \geq 1 - \epsilon$. The assumption about \bar{L} in the statement of the theorem also implies the existence of an open set \bar{B}_ϵ such that $\int_{\bar{B}_\epsilon} d\bar{P} > 1 - \epsilon$ and $\bar{L} : \bar{B}_\epsilon \mapsto \mathbb{R}$ is Lipschitz, since \bar{P} is absolutely continuous with respect to P_0 (cf. Example 3 on page 55 of Pollard (2002)). With $B = B_\epsilon \cap \bar{B}_\epsilon$, $\bar{L}^i : B \mapsto \mathbb{R}$ with $\bar{L}^i = 1/\bar{L}$ is thus bounded and Lipschitz. Furthermore, since $S \setminus B$ is closed, there exists a Lipschitz function $\ell : S \mapsto [0, 1]$ that is zero on $S \setminus B$ and for which $\int \ell d\bar{P} \geq 1 - 3\epsilon$ (see Pollard (2002), p. 172-173 for an explicit construction). For future reference, define \mathcal{B} and \mathcal{B}_ϵ to be the indicator functions of B and B_ϵ , respectively, and note that $\mathcal{B}\ell = \mathcal{B}_\epsilon\ell = \ell$. Define the scalar sequence

$$\begin{aligned} \kappa_T &= \int (\ell \circ h_T) dF_{0,T} / \int (\ell \circ h_T) (\bar{L}^i \circ h_T) d\bar{F}_T \\ &= \int \ell dP_{0,T} / \int \ell \bar{L}^i d\bar{P}_T \end{aligned}$$

and note that $\kappa_T \rightarrow 1$ because $\int \ell dP_{0,T} \rightarrow \int \ell dP_0$ and $\int \ell \bar{L}^i d\bar{P}_T \rightarrow \int \ell \bar{L}^i d\bar{P} = \int \mathcal{B}\ell \bar{L}^i \bar{L} dP_0 = \int \ell dP_0$ by the Continuous Mapping Theorem. Further, define the probability distribution G_T on $\mathfrak{B}(\mathbb{R}^{nT})$ via

$$\int_A dG_T = \kappa_T \int_A (\ell \circ h_T) (\bar{L}^i \circ h_T) d\bar{F}_T + \int_A ((1 - \ell) \circ h_T) dF_{0,T}$$

for any $A \in \mathfrak{B}(\mathbb{R}^{nT})$. Then by construction, h_T induces the probability distribution Q_T under G_T , where Q_T satisfies

$$\int_A dQ_T = \kappa_T \int_A \ell \bar{L}^i d\bar{P}_T + \int_A (1 - \ell) dP_{0,T}$$

for any $A \in \mathfrak{B}(S)$. Now

$$\begin{aligned} \Delta_{BL}(Q_T, P_0) &= \sup_{\|\vartheta\|_{BL} \leq 1} \left| \int \vartheta(dQ_T - dP_0) \right| \\ &\leq \sup_{\|\vartheta\|_{BL} \leq 1} \left| \int \vartheta \ell \bar{L}^i (\kappa_T d\bar{P}_T - d\bar{P}) \right| + \sup_{\|\vartheta\|_{BL} \leq 1} \left| \int \vartheta(1 - \ell)(dP_{0,T} - dP_0) \right| \\ &\leq \|\ell \bar{L}^i\|_{BL} (\Delta_{BL}(\bar{P}_T, \bar{P}) + |\kappa_T - 1|) + \|1 - \ell\|_{BL} \Delta_{BL}(P_{0,T}, P_0) \end{aligned}$$

where the manipulations after the first inequality use $\int \vartheta \ell dP_0 = \int \vartheta \ell \bar{L}^i d\bar{P}$ and the second inequality exploits that $\|\cdot\|_{BL}$ is a submultiplicative norm on the set of bounded Lipschitz

functions $S \mapsto \mathbb{R}$ (cf. Proposition 11.2.1 of Dudley (2002)). Also,

$$\begin{aligned} |\kappa_T - 1| &= \left| \frac{\int \ell dP_{0,T}}{\int \ell \bar{L}^i d\bar{P}_T} - \frac{\int \ell dP_0}{\int \ell \bar{L}^i d\bar{P}} \right| \\ &\leq \frac{\|\ell \bar{L}^i\|_{BL} \Delta_{BL}(\bar{P}_T, \bar{P}) + \|\ell\|_{BL} \Delta_{BL}(P_{0,T}, P_0)}{\int \ell \bar{L}^i d\bar{P}_T}. \end{aligned}$$

Thus, $\lim_{T \rightarrow \infty} \Delta_{BL}(Q_T, P_0)/\delta_T = 0$, so that for large enough T , (18) implies $\limsup_{T \rightarrow \infty} \int \varphi_T dG_T \leq \alpha$.

Now define the probability measures \tilde{F}_T via

$$\begin{aligned} \int_A d\tilde{F}_T &= \tilde{\kappa}_T \int_A (\mathcal{B}_\epsilon \circ h_T)(\bar{L} \circ h_T) dG_T \\ &= \tilde{\kappa}_T \kappa_T \int_A (\ell \circ h_T) d\bar{F}_T + \tilde{\kappa}_T \int_A (\mathcal{B}_\epsilon(1 - \ell) \bar{L} \circ h_T) dF_{0,T} \end{aligned}$$

for any $A \in \mathfrak{B}(\mathbb{R}^{nT})$, where $\tilde{\kappa}_T = 1/(\kappa_T \int (\ell \circ h_T) d\bar{F}_T + \int (\mathcal{B}_\epsilon(1 - \ell) \bar{L} \circ h_T) dF_{0,T}) \rightarrow \tilde{\kappa} = 1/\int (\ell + \mathcal{B}_\epsilon - \mathcal{B}_\epsilon \ell) d\bar{P} = 1/\int \mathcal{B}_\epsilon d\bar{P}$ and $1 \leq 1/\int \mathcal{B}_\epsilon d\bar{P} \leq 1 + 2\epsilon$. By the Neyman-Pearson Lemma, the best test of $\tilde{H}_0 : Y_T \sim G_T$ against $\tilde{H}_1 : Y_T \sim \tilde{F}_T$ thus rejects for large values of $(\mathcal{B}_\epsilon \bar{L}) \circ h_T$, i.e. $\bar{L} \circ h_T$.

For any T , denote by $\tilde{\varphi}_T^* : \mathbb{R}^{nT} \mapsto [0, 1]$ the test that rejects for large values of $\bar{L} \circ h_T$ of level $\int \tilde{\varphi}_T^* dG_T = \max(\int \varphi_T dG_T, \alpha)$, so that $\int (\tilde{\varphi}_T^* - \varphi_T) d\tilde{F}_T \geq 0$ for all T . By LeCam's First Lemma (Lemma 6.4 in van der Vaart (1998)), \tilde{F}_T is contiguous to G_T , since under G_T , the Radon-Nikodym derivative $\tilde{\kappa}_T(\mathcal{B}_\epsilon \bar{L}) \circ h_T$ converges weakly to the distribution $\tilde{\kappa} \mathcal{B}_\epsilon \bar{L} P_0$ by the Continuous Mapping Theorem, and $\int \tilde{\kappa} \mathcal{B}_\epsilon \bar{L} dP_0 = 1$. Since both $\tilde{\varphi}_T^*$ and $\hat{\varphi}_T^*$ reject for large values of $\bar{L} \circ h_T$ and are of asymptotic level α , we have $\int |\tilde{\varphi}_T^* - \hat{\varphi}_T^*| dG_T \rightarrow 0$, so that by contiguity, also $\int |\tilde{\varphi}_T^* - \hat{\varphi}_T^*| d\tilde{F}_T \rightarrow 0$. Thus $\limsup_{T \rightarrow \infty} \int (\varphi_T - \hat{\varphi}_T^*) d\tilde{F}_T \leq 0$. To complete the proof, note that the total variation distance between \tilde{F}_T and \bar{F}_T is bounded above by

$$\begin{aligned} \int |1 - \tilde{\kappa}_T \kappa_T (\ell \circ h_T)| d\bar{F}_T &\leq |1 - \tilde{\kappa}_T \kappa_T| + |\tilde{\kappa}_T \kappa_T| \int (1 - \ell) d\bar{P}_T \\ &\rightarrow 1/\int \mathcal{B}_\epsilon d\bar{P} - 1 + (1 - \int \ell d\bar{P})/\int \mathcal{B}_\epsilon d\bar{P} \leq 8\epsilon \end{aligned}$$

so that $\limsup_{T \rightarrow \infty} \int (\varphi_T - \hat{\varphi}_T^*) d\bar{F}_T \leq 8\epsilon$, and the result follows, since the choice of $1/2 > \epsilon > 0$ was arbitrary.

References

- ANDERSON, T. W., AND H. RUBIN (1949): “Estimators of the Parameters of a Single Equation in a Complete Set of Stochastic Equations,” *The Annals of Mathematical Statistics*, 21, 570–582.
- ANDREWS, D. W. K. (1993): “Tests for Parameter Instability and Structural Change with Unknown Change Point,” *Econometrica*, 61, 821–856.
- 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.
- (2008): “Efficient Two-Sided Nonsimilar Invariant Tests in IV Regression with Weak Instruments,” *Journal of Econometrics*, 146, 241–254.
- ANDREWS, D. W. K., AND W. PLOBERGER (1994): “Optimal Tests When a Nuisance Parameter Is Present Only under the Alternative,” *Econometrica*, 62, 1383–1414.
- BICKEL, P. J., C. A. J. KLAASSEN, Y. RITOV, AND J. A. WELLNER (1998): *Efficient and Adaptive Estimation for Semiparametric Models*. Springer-Verlag, New York.
- BILLINGSLEY, P. (1968): *Convergence of Probability Measure*. Wiley, New York.
- BREITUNG, J. (2002): “Nonparametric Tests for Unit Roots and Cointegration,” *Journal of Econometrics*, 108, 343–363.
- CHOI, A., W. J. HALL, AND A. SCHICK (1996): “Asymptotically Uniformly Most Powerful Tests in Parametric and Semiparametric Models,” *Annals of Statistics*, 24, 841–861.
- DAVIDSON, J. (2002): “Establishing Conditions for the Functional Central Limit Theorem in Nonlinear and Semiparametric Time Series Processes,” *Journal of Econometrics*, 106, 243–269.
- (2007): “When is a Time Series $I(0)$?,” *University of Exeter Working Paper*.
- DE JONG, R. M., AND J. DAVIDSON (2000): “The Functional Central Limit Theorem and Weak Convergence to Stochastic Integrals I: Weakly Dependent Processes,” *Econometric Theory*, 16, 621–642.

- DUDLEY, R. M. (2002): *Real Analysis and Probability*. Cambridge University Press, Cambridge, UK.
- ELLIOTT, G. (1999): “Efficient Tests for a Unit Root When the Initial Observation is Drawn From its Unconditional Distribution,” *International Economic Review*, 40, 767–783.
- ELLIOTT, G., AND M. JANSSON (2003): “Testing for Unit Roots with Stationary Covariates,” *Journal of Econometrics*, 115, 75–89.
- ELLIOTT, G., M. JANSSON, AND E. PESAVENTO (2005): “Optimal Power for Testing Potential Cointegrating Vectors with Known Parameters for Nonstationarity,” *Journal of Business and Economics Statistics*, 23, 34–48.
- ELLIOTT, G., AND U. K. MÜLLER (2006): “Efficient Tests for General Persistent Time Variation in Regression Coefficients,” *Review of Economic Studies*, 73, 907–940.
- ELLIOTT, G., T. J. ROTHENBERG, AND J. H. STOCK (1996): “Efficient Tests for an Autoregressive Unit Root,” *Econometrica*, 64, 813–836.
- HANSEN, B. E. (1990): “Convergence to a Stochastic Integrals for Dependent Heterogeneous Processes,” *Econometric Theory*, 8, 489–500.
- IBRAGIMOV, R., AND U. K. MÜLLER (2007): “T-Statistic Based Correlation and Heterogeneity Robust Inference,” *Harvard Institute of Economic Research Discussion Paper No. 2129*.
- JANSSON, M. (2005): “Point Optimal Tests of the Null Hypothesis of Cointegration,” *Journal of Econometrics*, 124, 187–201.
- (2008): “Semiparametric Power Envelopes for Tests of the Unit Root Hypothesis,” *Econometrica*, 76, 1103–1142.
- JANSSON, M., AND M. J. MOREIRA (2006): “Optimal Inference in Regression Models with Nearly Integrated Regressors,” *Econometrica*, 74, 681–714.
- LEHMANN, E. L., AND J. P. ROMANO (2005): *Testing Statistical Hypotheses*. Springer, New York.

- LI, H., AND U. K. MÜLLER (2007): “Valid Inference in Partially Unstable GMM Models,” *forthcoming in the Review of Economic Studies*.
- MCLEISH, D. L. (1974): “Dependent Central Limit Theorems and Invariance Principles,” *The Annals of Probability*, 2, 620–628.
- MOREIRA, M. J. (2001): “Tests with Correct Size When Instruments Can Be Arbitrarily Weak,” *Working Paper, Berkeley*.
- MÜLLER, U. K. (2007): “Comment on: Unit Root Testing in Practice Dealing with Uncertainty Over the Trend and Initial Condition,” *forthcoming in Econometric Theory*.
- (2008): “The Impossibility of Consistent Discrimination Between $I(0)$ and $I(1)$ Processes,” *Econometric Theory*, 24, 616–630.
- MÜLLER, U. K., AND G. ELLIOTT (2003): “Tests for Unit Roots and the Initial Condition,” *Econometrica*, 71, 1269–1286.
- MÜLLER, U. K., AND M. W. WATSON (2008a): “Low-Frequency Robust Cointegration Testing,” *Working paper, Princeton University*.
- (2008b): “Testing Models of Low-Frequency Variability,” *Econometrica*, 76, 979–1016.
- NEWBY, W. K. (1990): “Semiparametric Efficiency Bounds,” *Journal of Applied Econometrics*, 99–135, 5.
- NYBLOM, J. (1989): “Testing for the Constancy of Parameters Over Time,” *Journal of the American Statistical Association*, 84, 223–230.
- PHILLIPS, P. C. B. (1988): “Weak Convergence of Sample Covariance Matrices to Stochastic Integrals Via Martingale Approximations,” *Econometric Theory*, 4, 528–533.
- PHILLIPS, P. C. B., AND V. SOLO (1992): “Asymptotics for Linear Processes,” *Annals of Statistics*, 20, 971–1001.
- POLLARD, D. (2001): “Contiguity,” Working paper, Yale University.

- (2002): *A User's Guide to Measure Theoretic Probability*. Cambridge University Press, Cambridge, UK.
- ROTHENBERG, T. J., AND J. H. STOCK (1997): "Inference in a Nearly Integrated Autoregressive Model with Nonnormal Innovations," *Journal of Econometrics*, 80, 269–286.
- SCHERVISH, M. J. (1995): *Theory of Statistics*. Springer, New York.
- SHIRYAEV, A. N., AND V. G. SPOKOINY (2000): *Statistical Experiments and Decisions: Asymptotic Theory*. World Scientific Publishing, Singapore.
- SOWELL, F. (1996): "Optimal Tests for Parameter Instability in the Generalized Method of Moments Framework," *Econometrica*, 64, 1085–1107.
- STAIGER, D., AND J. H. STOCK (1997): "Instrumental Variables Regression with Weak Instruments," *Econometrica*, 65, 557–586.
- STEIN, C. (1956): "Efficient Nonparametric Estimation and Testing," in *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability, Volume*, pp. 187–196, Berkeley. University of California Press.
- STOCK, J. H. (1994): "Deciding Between $I(1)$ and $I(0)$," *Journal of Econometrics*, 63, 105–131.
- STOCK, J. H., AND M. W. WATSON (1996): "Confidence Sets in Regressions with Highly Serially Correlated Regressors," *Working Paper, Harvard University*.
- STOCK, J. H., AND J. H. WRIGHT (2000): "GMM with Weak Identification," *Econometrica*, 68, 1055–1096.
- VAN DER VAART, A. W. (1998): *Asymptotic Statistics*. Cambridge University Press, Cambridge, UK.
- WHITE, H. (2001): *Asymptotic Theory for Econometricians*. Academic Press, San Diego, revised edn.
- WOOLDRIDGE, J. M., AND H. WHITE (1988): "Some Invariance Principles and Central Limit Theorems for Dependent Heterogeneous Processes," *Econometric Theory*, 4, 210–230.