

HOW TO AVOID THE ZERO-POWER TRAP IN TESTING FOR CORRELATION*

DAVID PREINERSTORFER
Université libre de Bruxelles

In testing for correlation of the errors in regression models, the power of tests can be very low for strongly correlated errors. This counterintuitive phenomenon has become known as the “zero-power trap.” Despite a considerable amount of literature devoted to this problem, mainly focusing on its detection, a convincing solution has not yet been found. In this article, we first discuss theoretical results concerning the occurrence of the zero-power trap phenomenon. Then, we suggest and compare three ways to avoid it. Given an initial test that suffers from the zero-power trap, the method we recommend for practice leads to a modified test whose power converges to 1 as the correlation gets very strong. Furthermore, the modified test has approximately the same power function as the initial test and thus approximately preserves all of its optimality properties. We also provide some numerical illustrations in the context of testing for network generated correlation.

1. INTRODUCTION

Testing whether the errors in a regression model are uncorrelated is a standard problem in econometrics. For many forms of correlation under the alternative, there are well-established tests available. Two prominent examples are the Durbin–Watson test for serial autocorrelation and the Cliff–Ord test for spatial autocorrelation. Nevertheless, this type of testing problem is not completely solved, not even in the Gaussian case. This is partly due to the fact that tests for correlation, including the well-established tests mentioned before, do not always behave as they ideally should in finite samples: Whereas the size of most tests can be easily controlled, at least under suitable distributional assumptions, such as Gaussianity, their power function can attain very small values in regions of the alternative where the correlation is very strong. This, however, does not match with the intuition that strong correlations should be *easily* detectable from the data, i.e., that the power of a test for correlation should be close to 1 if the degree of correlation in the errors is very strong.

*I thank the Editor, Co-Editor, and two referees for helpful comments. Financial support by the Program of Concerted Research Actions (ARCs) of the Université libre de Bruxelles is gratefully acknowledged. Address correspondence to David Preinerstorfer, European Center for Advanced Research in Economics and Statistics and Solvay Brussels School of Economics and Management, Université libre de Bruxelles, Avenue Franklin Roosevelt 50, 1050 Bruxelles, Belgium; e-mail: david.preinerstorfer@ulb.be.

Overlooking strong correlations in the errors can have severe practical consequences. For example, Kelly (2019) discusses 27 prominent examples from the economics literature, where strong spatial autocorrelation in the errors potentially led to overly large t -statistics and spurious conclusions. Among others, the author recommends supplementing the analysis by a test statistic for spatial autocorrelation (note that the author uses the term Moran's I statistic which we refer to as the Cliff–Ord test statistic). This suggestion anticipates that such tests are automatically sensitive to strong correlations, which, however, is not guaranteed.

That the power function of a test for correlation can drop to 0 as the correlation increases was first formally established in Krämer (1985), who considered the power function of the Durbin–Watson test in testing for serial autocorrelation. The results in Krämer (1985) were extended in later work by Zeisel (1989), Krämer and Zeisel (1990), and Löbus and Ritter (2000). Kleiber and Krämer (2005) obtained similar results for the Durbin–Watson test when the disturbances are fractionally integrated. Krämer (2005) proved related results for Cliff–Ord-type tests in cases where the regression errors are spatially autocorrelated. A unifying general theory that neither relies on the specific form of correlation nor on very special structural properties of the tests was developed recently in Martellosio (2010) and Preinerstorfer and Pötscher (2017). We refer the interested reader to the latter articles for formal results and a thorough discussion of the literature.

The major practical value of the just mentioned articles is of a diagnostic nature: They provide conditions which depend on observable quantities only and which let a user detect whether a particular test is subject to the *zero-power trap*, i.e., whether its power function drops to 0 as the correlation increases. This is important because if it turns out that an initial test is subject to this trap one may want to use another test. However, one is then confronted with the problem of finding a test that avoids the zero-power trap. One complication is as follows: Typically, the initial test was chosen for a reason, i.e., for its “optimal” power properties in certain regions of the parameter space (think of a locally best invariant test). In such situations, one would not just like to use *some* other test that avoids the zero-power trap. Much more likely, one would prefer to *slightly modify* the initial test in such a way that its optimality properties are preserved, at least approximately, but such that its modified version does not suffer from the zero-power trap. Compared with the amount of literature that concentrates on deriving diagnostic tools for detecting the zero-power trap, the attention that has been paid to the question *how* one can construct tests which do not suffer from the zero-power trap is much less. Furthermore, it is not clear how to obtain said “optimality-preserving” modifications. The main contribution of the present article is to fill this gap. In the following paragraphs, we provide an overview of the article's structure together with a more detailed summary of our contributions.

In Section 2, we introduce the framework: the model and the testing problem, some notational conventions, and an important class of tests. In Section 3, we formally define the zero-power trap phenomenon, obtain some sufficient conditions for it from results in Preinerstorfer and Pötscher (2017), and then

consider in our general framework the question how often, i.e., for “how many” design matrices, the zero-power trap actually arises. We answer this question in Propositions 3.4 and 3.6. The former proposition proves (and generalizes) an observation already made in the discussion section of Krämer (1985). The latter proposition is obtained by generalizing an argument in Martellosio (2012), who considered the same question in a spatial autoregressive setting. Essentially, these two propositions show (for the tests based on the specific family of test statistics and the corresponding critical values considered), respectively, that (i) the zero-power trap arises for generic design matrices (i.e., up to a Lebesgue null set of exceptional matrices) for small enough critical values; and (ii) for any critical value that leads to a size in $(0, 1)$, there exists an open set of design matrices for which the zero-power trap arises.

In Section 4, we present three ways to avoid the zero-power trap: In Section 4.1, we briefly discuss a test for which Preinerstorfer and Pötscher (2017) have shown that it does not suffer from the zero-power trap. This test typically does not have very favorable power properties, apart from the fact that it avoids the zero-power trap. We shall mainly use it later as a building block in our construction of “optimality-preserving” tests. In Section 4.2, we discuss tests that incorporate artificial regressors to avoid the zero-power trap. The suggestion of adding artificial regressors to the regression and to use “optimal” tests in this expanded model is present already in Krämer (1985), who observed numerically that adding the intercept to a regression without intercept helps to avoid the zero-power trap for the Durbin–Watson test. Our theoretical results in Section 4.2 exploit results in Preinerstorfer and Pötscher (2017) and are related to the methods in Preinerstorfer and Pötscher (2016) and Preinerstorfer (2017), who considered the construction of tests with good size and power properties for testing restrictions on the regression coefficient vector. While the tests in Section 4.2 are “optimality-preserving” to some extent (more specifically, they often have the same optimality property as initial tests, but within a smaller class of tests), it turns out that this solution to the zero-power trap is not ideal. For example, the power function of these tests does not increase to 1, as the strength of the correlation increases (which is the case for the approach outlined in Section 4.1).

In Section 4.3, we construct optimality-preserving modifications avoiding the zero-power-trap out of an initial test that suffers from the zero-power trap. Our approach overcomes the limitations of the approaches discussed in Sections 4.1 and 4.2. In particular, our method leads to tests that have approximately the same power properties as the initial test. Furthermore, their power converges to 1 as the strength of the correlation increases. The construction is inspired by the power enhancement principle of Fan et al. (2015) in the formulation used in Section 3 of Kock and Preinerstorfer (2019). The basic idea of this principle is to improve the asymptotic power of an initial test by using another test, a power enhancement component, which has better asymptotic power properties than the initial test in certain regions of the alternative. Since the theory in Fan et al. (2015) and Kock and Preinerstorfer (2019) is asymptotic, and *the present article is concerned exclusively*

with finite sample properties, their results do not apply here. Nevertheless, we can adapt the underlying heuristic to our context: Given an initial test that suffers from the zero-power trap, but has favorable power properties in other regions of the alternative, we “combine” this initial test with the test from Section 4.1 to obtain an “enhanced” test.

In Section 5, we compare the approaches for avoiding the zero-power trap discussed in Section 4 numerically. We reconsider an example in Krämer (2005) in which the Cliff–Ord test turns out to suffer from the zero-power trap. Section 6 concludes. All proofs are collected in Appendixes A–C.

2. FRAMEWORK

In the present section, we introduce the model, the testing problem, some notation, and we discuss an important class of tests. Most of the notational conventions and terminology we use are standard and coincide to a large extent with the ones in Preinerstorfer and Pötscher (2017). We repeat them here for the convenience of the reader.

2.1. Model and Testing Problem

We consider the linear model

$$\mathbf{y} = X\boldsymbol{\beta} + \mathbf{u}, \quad (1)$$

where $X \in \mathbb{R}^{n \times k}$ is a nonstochastic matrix of rank k , with $0 < k < n$, and where $\boldsymbol{\beta} \in \mathbb{R}^k$ is the regression coefficient vector. The disturbance vector \mathbf{u} is assumed to be Gaussian with mean 0 and covariance matrix $\sigma^2 \Sigma(\rho)$. Here, $\Sigma(\cdot)$ is a known function from $[0, a)$ to the set of symmetric and positive definite $n \times n$ matrices, and a is a prespecified positive real number. Without loss of generality, we assume throughout that $\Sigma(0)$ equals the identity matrix I_n . The parameters $\boldsymbol{\beta} \in \mathbb{R}^k$, $\sigma \in (0, \infty)$, and $\rho \in [0, a)$ are unknown.

Remark 2.1. In the context of the zero-power trap, the upper endpoint a of the parameter space of ρ will always be such that the covariance matrix $\Sigma(\rho)$ is “close to” being rank deficient for ρ close to a . This will be made precise in Assumption 1 below and the discussion preceding it. The parameter ρ being close to a may thus be interpreted as \mathbf{u} being strongly correlated. The concrete value of a depends on the specific parameterization used and does not have any intrinsic meaning, as one can in principle always enforce $a = 1$ by a suitable reparameterization. In a spatial autoregressive model in the usual parameterization, for example, a corresponds to the inverse of the eigenvalue with largest absolute value of the weights matrix (cf. Section 5 for more details).

The Gaussianity assumption could be relaxed considerably. It is imposed mainly to avoid technical conditions that do not deliver deeper insights into the problem. For example, we could replace the Gaussianity assumption by the assumption that

the distribution of the error vector \mathbf{u} is elliptically symmetric without changing any of our results. This and other generalizations are discussed in detail in Section 3 of Preinerstorfer and Pötscher (2017).

Denoting the Gaussian probability measure with mean $X\beta$ and covariance matrix $\sigma^2\Sigma(\rho)$ by $P_{\beta,\sigma,\rho}$, we see that the model (1) induces the parametric family of distributions

$$\{P_{\beta,\sigma,\rho} : \beta \in \mathbb{R}^k, \sigma \in (0, \infty), \rho \in [0, a)\} \quad (2)$$

on the sample space \mathbb{R}^n equipped with its Borel σ -algebra. The expectation operator with respect to (w.r.t.) $P_{\beta,\sigma,\rho}$ will be denoted by $\mathbb{E}_{\beta,\sigma,\rho}$. Note that the set of probability measures in the previous display is dominated by Lebesgue measure $\mu_{\mathbb{R}^n}$ on the Borel sets of \mathbb{R}^n , because $\Sigma(\rho)$ is positive definite, for every $\rho \in [0, a)$, by assumption.

In the family of distributions (2), we are interested in the testing problem $\rho = 0$ against $\rho > 0$. More precisely, the testing problem is

$$H_0 : \rho = 0, \beta \in \mathbb{R}^k, 0 < \sigma < \infty \quad \text{against} \quad H_1 : \rho > 0, \beta \in \mathbb{R}^k, 0 < \sigma < \infty, \quad (3)$$

with the implicit understanding that always $\rho \in [0, a)$. In this testing problem, the parameter ρ is the target of inference, and the regression coefficient vector β and the parameter σ are nuisance parameters.

Two specific examples that received a considerable amount of attention in the econometrics literature and which fit into the above framework are testing for positive serial autocorrelation and testing for spatial autocorrelation, cf. Examples 2.1 and 2.2 in Preinerstorfer and Pötscher (2017) for details and a discussion of related literature. See also Section 5 below for more information on testing for spatial autocorrelation and related numerical results.

2.2. Notation, Invariance, and an Important Class of Tests

2.2.1. *Notation.* All matrices we shall consider are real matrices, the transpose of a matrix A is denoted by A' , and the linear space spanned by the columns of A is denoted by $\text{span}(A)$. Given a linear subspace L of \mathbb{R}^n , the symbol Π_L denotes the orthogonal projection onto L , and L^\perp denotes the orthogonal complement of L . Given an $n \times m$ matrix Z of rank m with $0 \leq m < n$, we denote by C_Z a matrix in $\mathbb{R}^{(n-m) \times n}$ such that $C_Z C_Z' = I_{n-m}$ and $C_Z' C_Z = \Pi_{\text{span}(Z)^\perp}$, where I_r denotes the identity matrix of dimension r . We observe that every matrix whose rows form an orthonormal basis of $\text{span}(Z)^\perp$ satisfies these two conditions and vice versa. Hence, any two choices for C_Z are related by premultiplication by an orthogonal matrix. Let l be a positive integer. If A is an $l \times l$ matrix and $\lambda \in \mathbb{R}$ is an eigenvalue of A , we denote the corresponding eigenspace by $\text{Eig}(A, \lambda)$. The eigenvalues of a symmetric matrix $B \in \mathbb{R}^{l \times l}$ ordered from smallest to largest and counted with their multiplicities are denoted by $\lambda_1(B), \dots, \lambda_l(B)$. We shall sometimes denote $\lambda_1(B)$ by $\lambda_{\min}(B)$, and $\lambda_l(B)$ by $\lambda_{\max}(B)$. Lebesgue measure on the Borel σ -algebra

of $\mathbb{R}^{n \times l}$ shall be denoted by $\mu_{\mathbb{R}^{n \times l}}$, and \Pr is used as a generic symbol for a probability measure. The Euclidean norm of a vector is denoted by $\|\cdot\|$, a symbol that is also used to denote a matrix norm.

2.2.2. Invariance, an Important Class of Tests, and Size-Controlling Critical Values. Given a matrix $Z \in \mathbb{R}^{n \times m}$ with column rank m and where $1 \leq m < n$, define the group of bijective transformations (the group action being composition of functions)

$$G_Z := \{g_{\gamma, \theta} : \gamma \in \mathbb{R} \setminus \{0\}, \theta \in \mathbb{R}^m\},$$

where $g_{\gamma, \theta} : \mathbb{R}^n \rightarrow \mathbb{R}^n$ denotes the function $y \mapsto \gamma y + Z\theta$.

Under our distributional assumptions (and if additionally all parameters of the model are identifiable), the testing problem in equation (3) is invariant w.r.t. the group G_X (cf. Section 6 in Lehmann and Romano (2005)). It thus appears reasonable to consider tests that are G_X -invariant, a property shared by most commonly used tests. Recall that a function f defined on the sample space (e.g., a test or a test statistic) is called invariant w.r.t. G_X if and only if, for every $y \in \mathbb{R}^n$ and every $g_{\gamma, \theta} \in G_X$, it holds that $f(y) = f(g_{\gamma, \theta}(y))$. A subset A of \mathbb{R}^n will be called invariant w.r.t. G_X if the indicator function $\mathbf{1}_A$ is G_X -invariant.

In addition to being G_X -invariant, most tests for (3) used in practice are nonrandomized, i.e., they are indicator functions of Borel sets (their corresponding rejection regions). An important class of such tests is based on rejection regions of the form

$$\Phi_{B,c} = \Phi_{B,C_X,c} = \{y \in \mathbb{R}^n : T_B(y) > c\},$$

where $c \in \mathbb{R}$ is a critical value and the test statistic

$$T_B(y) = T_{B,C_X}(y) = \begin{cases} y' C'_X B C_X y / \|C_X y\|^2 & \text{if } y \notin \text{span}(X) \\ \lambda_1(B) & \text{if } y \in \text{span}(X). \end{cases} \tag{4}$$

Here, $B \in \mathbb{R}^{(n-k) \times (n-k)}$ is a symmetric matrix, which typically depends on X and the function Σ . Recall that the matrix C_X satisfies $C_X C'_X = I_{n-k}$ and $C'_X C_X = \Pi_{\text{span}(X)^\perp}$ (cf. Section 2.2.1). Clearly, the test statistic T_B is G_X -invariant. Note furthermore that in case $\lambda_1(B) = \lambda_{n-k}(B)$, the test statistic T_B is constant everywhere on \mathbb{R}^n . Therefore, such a choice of B is uninteresting for practical purposes. Note also that assigning the value $\lambda_1(B)$ (instead of any other value) to the test statistic on $\text{span}(X)$ has no effect on rejection probabilities, because $P_{\beta, \sigma, \rho}$ is absolutely continuous w.r.t. $\mu_{\mathbb{R}^n}$, for every $\beta \in \mathbb{R}^k$, $\sigma \in (0, \infty)$, and $\rho \in [0, a)$, and because $\text{span}(X)$ being of dimension $k < n$ implies $\mu_{\mathbb{R}^n}(\text{span}(X)) = 0$.

The following remark discusses two particularly important choices of B .

Remark 2.2. Under regularity conditions and excluding degenerate cases, point-optimal invariant (w.r.t. G_X) tests and locally best invariant (w.r.t. G_X) tests for the testing problem (3) reject for large values of a test statistic T_B as in equation (4):

- (a) Point-optimal invariant tests against the alternative $\bar{\rho} \in (0, a)$ are obtained for $B = -(C_X \Sigma(\bar{\rho}) C_X')^{-1}$.
- (b) Locally best invariant tests are obtained for $B = C_X \dot{\Sigma}(0) C_X'$, for $\dot{\Sigma}(0)$ the derivative of Σ at $\rho = 0$, ensured to exist under the aforementioned regularity conditions, see, e.g., King and Hillier (1985).

Note that a test statistic T_B based on any of the two matrices B in the preceding enumeration does *not* depend on the specific choice of C_X , as any two choices of C_X differ only by premultiplication of an orthogonal matrix. However, for matrices B of a different form than (a) or (b), the test statistic T_B may also depend on the choice of C_X , a dependence which is typically suppressed in our notation.

The main focus of the present article concerns power properties of tests based on a test statistic as in (4) for the testing problem (3). Before investigating power properties of a test, one needs to ensure that its size does not exceed a given value of significance α . While this can be a nontrivial problem in general, achieving size control through the choice of a proper critical value turns out to be an easy task here. More specifically, the following lemma shows that exact size control for tests based on a test statistic T_B introduced in equation (4) is possible at all levels of significance in the leading case $\lambda_1(B) < \lambda_{n-k}(B)$. The subsequent remark discusses numerical aspects.

LEMMA 2.3. *Let $B \in \mathbb{R}^{(n-k) \times (n-k)}$ be symmetric and such that $\lambda_1(B) < \lambda_{n-k}(B)$. Then, there exists a (unique) function $\kappa : [0, 1] \rightarrow [\lambda_1(B), \lambda_{n-k}(B)]$, such that, for every $\alpha \in [0, 1]$,*

$$P_{\beta, \sigma, 0}(\Phi_{B, \kappa(\alpha)}) = \alpha \quad \text{for every } \beta \in \mathbb{R}^k \text{ and every } \sigma \in (0, \infty).$$

Furthermore, κ is a strictly decreasing and continuous bijection.

Remark 2.4. The rejection probabilities of a G_X -invariant test for (3) do not depend on the parameters β and σ (cf. Remark 2.3 in Preinerstorfer and Pötscher (2017)). As a consequence, the exact critical value $\kappa(\alpha)$ from Lemma 2.3 can easily be obtained numerically: To this end, one can exploit the well-known fact that, for every $c \in \mathbb{R}$, the rejection probability $P_{\beta, \sigma, 0}(\Phi_{B, c}) = P_{0, 1, 0}(\Phi_{B, c})$ can be rewritten as the probability that the quadratic form

$$\mathbf{G}' [B - cI_{n-k}] \mathbf{G} > 0,$$

where \mathbf{G} is an $(n - k)$ -variate Gaussian random vector with mean 0 and covariance matrix I_{n-k} . This probability can be determined efficiently through an application of standard algorithms, e.g., the algorithm by Davies (1980). The critical value $\kappa(\alpha)$ can then be obtained numerically by simply using a root-finding algorithm to determine the unique root $\kappa(\alpha)$ of $c \mapsto P_{0, 1, 0}(\Phi_{B, c}) - \alpha$ on $[\lambda_1(B), \lambda_{n-k}(B)]$.

3. THE ZERO-POWER TRAP IN TESTING FOR CORRELATION

3.1. Definition and Sufficient Conditions

In the sequel, a test $\varphi : \mathbb{R}^n \rightarrow [0, 1]$ (measurable) for testing problem (3) is said to be subject to (or suffer from) the zero-power trap, if there exist $\beta \in \mathbb{R}^k$ and $\sigma \in (0, \infty)$ such that

$$\liminf_{\rho \rightarrow a} \mathbb{E}_{\beta, \sigma, \rho}(\varphi) = 0; \tag{5}$$

that is, if the power function of φ can get arbitrarily close to 0 as the strength of the correlation in the data, measured in terms of ρ , increases. Recall from Remark 2.4 that if φ is G_X -invariant, which is the case for most tests considered in this article, then $\mathbb{E}_{\beta, \sigma, \rho}(\varphi)$ does not depend on β and σ . In this case, if equation (5) holds for some $\beta \in \mathbb{R}^k$ and some $\sigma \in (0, \infty)$, it holds for every $\beta \in \mathbb{R}^k$ and every $0 < \sigma < \infty$.

Trivially, a test that vanishes $\mu_{\mathbb{R}^n}$ -almost everywhere suffers from the zero-power trap regardless of the specific covariance structure $\Sigma(\cdot)$. As already hinted at in Remark 2.1, for the zero-power trap to occur for other tests, it is *necessary* that $\Sigma(\rho)$ is “close to” being rank-deficient for ρ close to a . To see this, note that if $\Sigma(\rho)$ converges to a positive definite matrix as $\rho \rightarrow a$ and if a test φ suffers from the zero-power trap, then $\varphi = 0$ must hold $\mu_{\mathbb{R}^n}$ -almost everywhere (as a consequence of Scheffé’s lemma). In case $\Sigma(\rho)$ does not converge to a positive definite matrix as $\rho \rightarrow a$, this is no longer true and the situation becomes interesting, as also other tests can suffer from the zero-power trap.

A set of *sufficient* conditions that allows one to conclude whether a test φ is subject to the zero-power trap was developed in Martellosio (2010) and Preinerstorfer and Pötscher (2017). The underlying effect leading to (5) described in the latter article is a concentration effect in the (rescaled) distribution $P_{\beta, \sigma, \rho}$ for ρ close to a , due to $\Sigma(\rho)$ being close to a rank-one matrix when ρ is close to a . More precisely, Preinerstorfer and Pötscher (2017) obtained their sufficient conditions under the following property of the function Σ (cf. also Assumption 1 in Preinerstorfer and Pötscher (2017) and the discussion there showing that this condition is weaker than the one previously used by Martellosio (2010)):

Assumption 1. $\lambda_n^{-1}(\Sigma(\rho))\Sigma(\rho) \rightarrow ee'$ as $\rho \rightarrow a$, for some $e \in \mathbb{R}^n$.

For the convenience of the reader and for later use, we shall now formally state two immediate consequences of results in Preinerstorfer and Pötscher (2017). They provide sufficient conditions for the zero-power trap under Assumption 1. Specializing Theorem 2.7 and Remark 2.8 in Preinerstorfer and Pötscher (2017), one obtains the following “high-level” result.

THEOREM 3.1. *Suppose Assumption 1 holds. Let φ be a G_X -invariant test that is continuous at e and satisfies $\varphi(e) = 0$, where e is the vector figuring in Assumption 1. Then,*

$$\lim_{\rho \rightarrow a} \mathbb{E}_{\beta, \sigma, \rho}(\varphi) = 0 \quad \text{for every } \beta \in \mathbb{R}^k \text{ and every } \sigma \in (0, \infty). \tag{6}$$

In particular, if $\varphi = \mathbf{1}_W$ holds for some G_X -invariant Borel set $W \subseteq \mathbb{R}^n$, then (6) holds if e is not in the closure of W .

For the test with rejection region $\Phi_{B,\kappa(\alpha)}$ as discussed in Section 2.2.2, and where $\kappa(\alpha)$ is defined through Lemma 2.3, one obtains the following result from Corollary 2.21 of Preinerstorfer and Pötscher (2017).

THEOREM 3.2. *Suppose Assumption 1 holds and $e \notin \text{span}(X)$, where e is the vector figuring in Assumption 1. Let $B \in \mathbb{R}^{(n-k) \times (n-k)}$ be symmetric and such that $\lambda_1(B) < \lambda_{n-k}(B)$. Then, for every $\alpha \in (0, 1)$ such that $T_B(e) < \kappa(\alpha)$, we have*

$$\lim_{\rho \rightarrow a} P_{\beta,\sigma,\rho}(\Phi_{B,\kappa(\alpha)}) = 0 \quad \text{for every } \beta \in \mathbb{R}^k \text{ and every } \sigma \in (0, \infty). \tag{7}$$

Note that the sufficient conditions for the zero-power trap phenomenon pointed out in Theorems 3.1 and 3.2 depend on observable quantities only, and that they are thus *checkable* by the user. Therefore, a researcher interested in testing problem (3) can use these conditions to check whether or not a given test suffers from the zero-power trap before actually using this test. In particular, one can decide *not* to use a test that suffers from the zero-power trap. Before addressing the question *how to avoid* the zero-power trap, which was raised already in the introduction, we briefly pay some attention to the following question: ‘‘How often’’ does the zero-power trap actually arise? More specifically, in the important class of tests $\Phi_{B,c}$ introduced in Section 2.2.2, and most notably the tests discussed in Remark 2.2, the following question arises: For ‘‘how many’’ design matrices X does the zero-power trap arise? Answering this question is the content of the next section.

3.2. For ‘‘How Many’’ Design Matrices Does the Zero-Power Trap Arise?

We shall focus on the class of tests with rejection regions $\Phi_{B(X),c}$ introduced in Section 2.2.2. Since the question in the section title depends on the design matrix X , which is otherwise held fixed in this article, we shall make the dependence of B on X explicit by writing $B(X)$. Furthermore, we shall also write $P_{\beta,\sigma,\rho}^X$ to emphasize its dependence on the design matrix X . Note that while X may vary in this section, the function $\Sigma(\cdot)$ and its domain $[0, a)$ are fixed, i.e., do not vary with X .

In our first attempt to answer the question under consideration, we shall use the following simple consequence of Lemma 2.3 and Theorem 3.2, which provides conditions on X under which equation (7) holds for all ‘‘small’’ levels α .

LEMMA 3.3. *Suppose Assumption 1 holds and let e denote the vector figuring in that assumption. Let B be a function from the set of full column rank $n \times k$ matrices to the set of symmetric $(n - k) \times (n - k)$ -dimensional matrices. If an $n \times k$ matrix X satisfies*

$$\text{rank}(X) = k \text{ and } C_X e \notin \text{Eig}(B(X), \lambda_{n-k}(B(X))), \tag{8}$$

then $\lambda_1(B(X)) < \lambda_{n-k}(B(X))$, $P_{0,1,0}^X(\Phi_{B(X),T_{B(X)}(e)}) > 0$, and equation (7) holds for every $\alpha \in (0, P_{0,1,0}^X(\Phi_{B(X),T_{B(X)}(e)}))$.

For a class of functions $X \mapsto B(X)$ that includes the ones discussed in Remark 2.2, we shall now show that condition (8) is generically satisfied, unless the matrix $B(X)$ has a very exceptional form. The result is established under a restriction concerning the eigenspace corresponding to the largest eigenvalue of $B(X)$.

PROPOSITION 3.4. *Suppose that $k < n - 1$ and that Assumption 1 holds. Let B be a function from the set of full column rank $n \times k$ matrices to the set of symmetric $(n - k) \times (n - k)$ -dimensional matrices. Let $M \in \mathbb{R}^{n \times n}$ be a symmetric matrix that cannot be written as $c_1 I_n + c_2 ee'$ for real numbers c_1, c_2 with $c_2 \geq 0$, where e is the vector figuring in Assumption 1. Suppose further that, for every $X \in \mathbb{R}^{n \times k}$ of full column rank, a $C_X \in \mathbb{R}^{(n-k) \times n}$ satisfying $C_X C_X' = I_{n-k}$ and $C_X' C_X = \Pi_{\text{span}(X)^\perp}$ can be chosen such that*

$$\text{Eig}(B(X), \lambda_{n-k}(B(X))) = \text{Eig}(C_X M C_X', \lambda_{n-k}(C_X M C_X')). \tag{9}$$

Then, up to a subset of a $\mu_{\mathbb{R}^{n \times k}}$ -null set of exceptional matrices, every $X \in \mathbb{R}^{n \times k}$ satisfies (8). An immediate consequence is as follows: Given $\alpha \in (0, 1)$ denote by $\mathcal{X}(\alpha; B) \subseteq \mathbb{R}^{n \times k}$ the set of all $X \in \mathbb{R}^{n \times k}$ of rank k such that $\lambda_1(B(X)) < \lambda_{n-k}(B(X))$ and such that

$$\lim_{\rho \rightarrow a} P_{\beta, \sigma, \rho}^X(\Phi_{B(X), C_X, \kappa(\alpha)}) = 0 \quad \text{for every } \beta \in \mathbb{R}^k \text{ and every } \sigma \in (0, \infty).$$

Then, $\mathcal{X}(\alpha_2; B) \subseteq \mathcal{X}(\alpha_1; B)$ holds for $0 < \alpha_1 \leq \alpha_2 < 1$, and for any sequence α_m in $(0, 1)$ converging to 0, the complement of $\bigcup_{m \in \mathbb{N}} \mathcal{X}(\alpha_m; B)$ is contained in a $\mu_{\mathbb{R}^{n \times k}}$ -null set.

Remark 3.5. Note that, for $B(X) = C_X \dot{\Sigma}(0) C_X'$, Condition (9) in Proposition 3.4 is trivially satisfied with $M = \dot{\Sigma}(0)$. For $B(X) = -(C_X \Sigma(\bar{\rho}) C_X')^{-1}$ and $\bar{\rho} \in (0, a)$, it is easy to see that Condition (9) is satisfied with $M = \Sigma(\bar{\rho})$. Therefore, if for any of these two specific choices the additional condition holds that the respective M cannot be written as $c_1 I_n + c_2 ee'$ for real numbers c_1, c_2 where $c_2 \geq 0$ is satisfied, then Proposition 3.4 applies.

Proposition 3.4 shows that tests based on $T_{B(X)}$ suffer from the zero-power trap for “most” design matrices X , at least for small choices of α . The discussion section of Krämer (1985) contains a corresponding statement (without proof) in a special case.

Choosing α small is not completely uncommon in practice: Due to the fact that testing for correlation is often just one part of the econometric analysis, the actual level α employed in this test can be quite small. One example is specification testing. Another example is the situation where tests for correlation are “inverted” to build a confidence interval for ρ , which is then used for a Bonferroni-type

construction of a data-dependent critical value of another test (cf. Leeb and Pötscher (2017) for further information concerning such critical values).

Nevertheless, the question remains as to how “large” the set $\mathcal{X}(\alpha; B)$ actually is for a fixed α , such as the conventional $\alpha = .05$ or $\alpha = .01$. For example, Proposition 3.4 does not tell us whether or not the set of design matrices $\mathcal{X}(.01; B)$ is empty. Similarly, one can ask if $\mathcal{X}(.01; B)$ contains an open set, or if it has positive $\mu_{\mathbb{R}^{n \times k}}$ measure. The latter questions have already been considered in detail in the main results of Martellosio (2012) for point-optimal invariant and locally best invariant tests in the important context of spatial autoregressive regression models. Adopting his proof strategy, we establish the following proposition. The argument requires a different assumption on B than the one used in Proposition 3.4. First, the condition used now concerns the eigenspace of $B(X)$ corresponding to its *smallest* eigenvalue (as opposed to the condition on the largest eigenvalue used in Proposition 3.4). Second, continuity conditions are imposed, which are required for limiting arguments in the proof. As discussed in Remark 3.7 below, the assumptions are again satisfied in the leading choices for B discussed in Remark 2.2.

PROPOSITION 3.6. *Suppose that $k < n - 1$ and that Assumption 1 holds. Let B be a function from the set of full column rank $n \times k$ matrices to the set of symmetric $(n - k) \times (n - k)$ -dimensional matrices. Suppose there exists a function F from the set of $(n - k) \times (n - k)$ matrices to itself, such that, for every $X \in \mathbb{R}^{n \times k}$ of full column rank, $B(X) = F(C_X M C_X')$ holds for a suitable choice of $C_X \in \mathbb{R}^{(n-k) \times n}$ satisfying $C_X C_X' = I_{n-k}$ and $C_X' C_X = \Pi_{\text{span}(X)^\perp}$, and for $M \in \mathbb{R}^{n \times n}$ a symmetric matrix that cannot be written as $c_1 I_n + c_2 e e'$ for real numbers c_1, c_2 where $c_2 \geq 0$. Here, e is the vector figuring in Assumption 1. Suppose further that F is continuous at every element A , say, of the closure of $\{C_X M C_X' : X \in \mathbb{R}^{n \times k}, \text{rank}(X) = k\} \subseteq \mathbb{R}^{(n-k) \times (n-k)}$, and that, for every such A , we have*

$$\text{Eig}(F(A), \lambda_1(F(A))) = \text{Eig}(A, \lambda_1(A)).$$

Define $\mathcal{X}(\alpha; B) \subseteq \mathbb{R}^{n \times k}$ as in Proposition 3.4. Then, the following holds:

1. $\mathcal{X}(\alpha; B) \neq \emptyset$ holds for every $\alpha \in (0, 1)$;
2. suppose that, for every $z \in \mathbb{R}^n$, the function $X \mapsto T_{B(X), C_X}(z)$ is continuous at every $X \in \mathbb{R}^{n \times k}$ of full column rank such that $z \notin \text{span}(X)$. Then, for every $\alpha \in (0, 1)$, the interior of $\mathcal{X}(\alpha; B)$ is nonempty (and thus has positive $\mu_{\mathbb{R}^{n \times k}}$ measure).

Remark 3.7. Similar to Remark 3.5, we note that Proposition 3.6 can be applied to $B(X) = C_X \check{\Sigma}(0) C_X'$ (with $M = \check{\Sigma}(0)$ and F the identity function), or to $B(X) = -(C_X \Sigma(\bar{\rho}) C_X')^{-1}$, where $\bar{\rho} \in (0, a)$, (with $M = \Sigma(\bar{\rho})$ and F the function $A \mapsto -A^{-1}$, noting that this function satisfies the continuity requirement as $\Sigma(\bar{\rho})$ is positive definite) provided that the corresponding M matrix is not of the exceptional form $c_1 I_n + c_2 e e'$, for $c_2 \geq 0$. It is not difficult to show that the continuity requirement in Part 2 of the proposition is satisfied for these two choices

of B . For $B(X) = C_X \hat{\Sigma}(0) C_X'$, this is trivial. For $B(X) = -(C_X \Sigma(\bar{\rho}) C_X')^{-1}$, where $\bar{\rho} \in (0, a)$, an argument is given in Appendix B. Hence, we can conclude that unless $\hat{\Sigma}(0)$ or $\Sigma(\bar{\rho})$, respectively, is of the form $c_1 I_n + c_2 ee'$ for some nonnegative c_2 , the test $\Phi_{B(X), \kappa(\alpha)}$ suffers from the zero-power trap for every $\alpha \in (0, 1)$ for every X in a nonempty open set of design matrices.

Remark 3.8. We emphasize that Propositions 3.4 and 3.6 do *not* apply in the special case where $M = c_1 I_n + c_2 ee'$ holds for real numbers c_1, c_2 where $c_2 \geq 0$. On the one hand, it is clear that in case $c_2 = 0$ a test as in these two propositions with $M = c_1 I_n$ trivially breaks down, as the corresponding test statistics are then constant. But on the other hand, as already observed (for the special case $c_1 = 0$ and $c_2 = 1$) in Preinerstorfer and Pötscher (2017) in the discussion preceding their Remark 2.27, using tests based on $M = c_1 I_n + c_2 ee'$, for a $c_2 > 0$, indeed presents an opportunity to avoid the zero-power trap. This will be discussed more formally in Section 4.1.

From the results in the present section, we learn that for tests that satisfy certain structural properties, the zero-power trap arises for generic design matrices for α small enough. Furthermore, for every α , there exists (under suitable assumptions) a nonempty open set of design matrices, every element of which suffers from the zero-power trap. We would like to emphasize, however, that these results *do not rule out* the possibility that, for a given X , the actual level α needed such that the zero-power trap arises can be low (far outside the commonly used range of levels), or that given α the open set of design matrices for which the zero-power trap occurs is “small.” Numerical results that illustrate the “practical severity” of the zero-power trap in spatial regression models are provided in Section 3 of Krämer (2005), in particular his Table 1 is very interesting in this context, and further discussion and examples can be found in Martellosio (2010) and Martellosio (2012). These results seem to suggest that the zero-power trap occurs frequently for commonly used levels of significance in case $n - k$ is “small,” i.e., in “high-dimensional” scenarios, whereas if $n - k$ is large, the zero-power trap does not appear that frequently. However, this also depends on the dependence structure.

4. AVOIDING THE ZERO-POWER TRAP

Having provided some context and motivation, we now discuss three ways to avoid the zero-power trap. In Section 4.1, we expand on the observation just made in Remark 3.8. The strategy discussed in Section 4.2 is based on an idea involving artificial regressors. The method we recommend, however, builds on Section 4.1 and is introduced in Section 4.3. Our suggestion tries to overcome suboptimality properties of the other methods. As discussed in the introduction, the idea underlying our approach can be interpreted as a finite sample variant of the power enhancement principle of Fan et al. (2015), cf. also Section 3 in Kock and Preinerstorfer (2019).

4.1. Tests Based on T_B with $B = C_X e e' C'_X$

As discussed in Remark 3.8, tests based on the test statistic T_B with $B = C_X e e' C'_X$ do not satisfy the assumptions underlying Propositions 3.4 and 3.6. Hence, these two propositions do not let us conclude anything concerning the question “how often” the zero-power trap occurs for such tests. It turns out that these tests do not suffer from the zero-power trap for any $\alpha \in (0, 1)$ in case the additional condition $e \notin \text{span}(X)$ holds (note that if $e \in \text{span}(X)$ holds, the test statistic T_B with $B = C_X e e' C'_X$ is useless as it equals 0 for every $y \in \mathbb{R}^n$). As pointed out in Remark 3.8, this was already noted in Preinerstorfer and Pötscher (2017). For later use in Section 4.3, we state a corresponding result (which is an immediate consequence of Part 1 of Proposition 2.26 in Preinerstorfer and Pötscher (2017) together with G_X -invariance of T_B and our Lemma 2.3).

THEOREM 4.1. *Suppose that $k < n - 1$, that Assumption 1 holds, and that $e \notin \text{span}(X)$, where e is the vector figuring in Assumption 1. Then, for every $\alpha \in (0, 1)$, every $\beta \in \mathbb{R}^k$, and every $\sigma \in (0, \infty)$,*

$$\lim_{\rho \rightarrow a} P_{\beta, \sigma, \rho}(\Phi_{C_X e e' C'_X, \kappa(\alpha)}) = 1.$$

From this result, we conclude that in case $e \notin \text{span}(X)$, and whenever a test φ with size α is subject to the zero-power trap, one can alternatively use the test with rejection region $\Phi_{C_X e e' C'_X, \kappa(\alpha)}$ instead, which does not suffer from the zero-power trap. Moreover, the power of the test $\Phi_{C_X e e' C'_X, \kappa(\alpha)}$ even increases to 1 as $\rho \rightarrow a$. This is a desirable property, as it matches the intuition that strong correlations should be easily detectable from the data.

While avoiding the zero-power trap problem, the test $\Phi_{C_X e e' C'_X, \kappa(\alpha)}$ suffers from one major disadvantage: The power function of $\Phi_{C_X e e' C'_X, \kappa(\alpha)}$ can be, and often will be, quite low for values $\rho \in (0, a)$ distant from a . If the initial test φ , which was dismissed because it is subject to the zero-power trap, was chosen because of its good power properties in this region of the alternative, the test $\Phi_{C_X e e' C'_X, \kappa(\alpha)}$ will then not constitute a convincing alternative. This is illustrated in the example discussed in Section 5. A method that tries to take optimality properties of the initial test into account, at least for the classes of tests discussed in Remark 2.2, is discussed next.

4.2. Tests Based on Artificial Regressors

The sufficient condition for the zero-power trap in Theorem 3.2 requires that the vector e from Assumption 1 is *not* an element of $\text{span}(X)$. While this of course does not *prove* that the zero-power trap does not arise if $e \in \text{span}(X)$, this indeed turns out to be the case under an additional assumption (cf. Corollary 2.22 in Preinerstorfer and Pötscher (2017)). In this section, we shall exploit this fact. The method of avoiding the zero-power trap we discuss in this section “enforces” the condition $e \in \text{span}(X)$. More specifically, it is based on adding the vector e from

Assumption 1 as an “artificial” regressor to the design matrix (if it is not already an element of $\text{span}(X)$), and from there constructing tests as if this artificially expanded design matrix was the true one. As discussed in the introduction, the idea underlying the construction in the present section can be traced back to Krämer (1985).

To formally describe the artificial regressor-based method in our general setting, consider a situation where a researcher initially wants to use the test $\Phi_{B,\kappa(\alpha)}$ as in Section 2.2.2 with $\lambda_1(B) < \lambda_{n-k}(B)$, but discovers (e.g., by checking the sufficient conditions in Theorem 3.2) that $\Phi_{B,\kappa(\alpha)}$ suffers from the zero-power trap. Suppose further that the initial test $\Phi_{B,\kappa(\alpha)}$ has certain optimality properties (cf. Remark 2.2). The researcher then may not want to completely sacrifice the optimality properties of the initial test, which prevents the researcher from using the test just discussed in Section 4.1. Assume further that $e \notin \text{span}(X)$.

The trick now is to work with the design matrix $\bar{X} = (X, e)$ in the construction of a test statistic, assuming that $k + 1 < n$. More precisely, let \bar{B} be a symmetric $(n - k - 1) \times (n - k - 1)$ matrix (cf. Remark 4.2 below) and define the adjusted test statistic

$$\bar{T}_{\bar{B}}(y) = \bar{T}_{\bar{B}, C_{\bar{X}}}(y) = \begin{cases} y' C_{\bar{X}}' \bar{B} C_{\bar{X}} y / \|C_{\bar{X}} y\|^2 & \text{if } y \notin \text{span}(\bar{X}) \\ \lambda_1(\bar{B}) & \text{if } y \in \text{span}(\bar{X}). \end{cases}$$

Under the additional assumption that $\lambda_1(\bar{B}) < \lambda_{n-k-1}(\bar{B})$, one obtains¹ from Lemma 2.3, for every $\alpha \in (0, 1)$, the existence and uniqueness of a critical value $\bar{\kappa}(\alpha) \in (\lambda_1(\bar{B}), \lambda_{n-k-1}(\bar{B}))$, say, such that, for every $\beta \in \mathbb{R}^k$ and every $\sigma \in (0, \infty)$, it holds that

$$P_{\beta, \sigma, 0}(\{y \in \mathbb{R}^n : \bar{T}_{\bar{B}}(y) > \bar{\kappa}(\alpha)\}) = \alpha. \tag{10}$$

Finally, define the rejection region

$$\bar{\Phi}_{\bar{B}, \bar{\kappa}(\alpha)} := \{y \in \mathbb{R}^n : \bar{T}_{\bar{B}}(y) > \bar{\kappa}(\alpha)\}. \tag{11}$$

Remark 4.2. We think about \bar{B} as an “updated version” of B , i.e., as the matrix, one would use if \bar{X} was the underlying design matrix. For example, if the initial matrix B equals $C_X \dot{\Sigma}(0) C_X'$, one could use $\bar{B} = C_{\bar{X}} \dot{\Sigma}(0) C_{\bar{X}}'$, or if the initial matrix $B = -(C_X \Sigma(\bar{\rho}) C_X')^{-1}$, one could use $\bar{B} = -(C_{\bar{X}} \Sigma(\bar{\rho}) C_{\bar{X}}')^{-1}$. Recall that the rejection region (11) based on these two versions of \bar{B} corresponds to locally best invariant tests and point-optimal invariant tests, respectively, in the model where the true design matrix is \bar{X} (cf. Remark 2.2).

We shall now prove that the test with rejection region (11) does not suffer from the zero-power trap. The following result requires an additional assumption on

¹To obtain this statement, one needs to apply Lemma 2.3 to model (1), but with design matrix \bar{X} instead of X . Note that this leads to an “enlarged” model that encompasses the true model as a submodel, and that the distributions satisfying the null hypothesis in the true model also satisfy the null hypothesis in the enlarged model.

$\Sigma(\cdot)$. This is Assumption 4 in Preinerstorfer and Pötscher (2017) to which we refer the reader for equivalent formulations, examples, and further discussion.

Assumption 2. There exists a function $c : [0, a) \rightarrow (0, \infty)$, a normalized vector $e \in \mathbb{R}^n$, and a square root $L_*(\cdot)$ of $\Sigma(\cdot)$ such that

$$\Lambda := \lim_{\rho \rightarrow a} c(\rho) \Pi_{\text{span}(e)^\perp} L_*(\rho)$$

exists in $\mathbb{R}^{n \times n}$ and such that the linear map Λ is injective when restricted to $\text{span}(e)^\perp$.

The main result concerning artificial regressor-based tests is as follows.

THEOREM 4.3. *Suppose Assumptions 1 and 2 are satisfied with the same vector e , that $e \notin \text{span}(X)$, and that $k < n - 1$. Suppose further that \bar{B} is a symmetric $(n - k - 1) \times (n - k - 1)$ matrix such that $\lambda_1(\bar{B}) < \lambda_{n-k-1}(\bar{B})$. Then, for every $\alpha \in (0, 1)$, every $\beta \in \mathbb{R}^k$, and every $\sigma \in (0, \infty)$, it holds that*

$$0 < \lim_{\rho \rightarrow a} P_{\beta, \sigma, \rho}(\bar{\Phi}_{\bar{B}, \bar{\kappa}(\alpha)}) = \Pr(\bar{T}_{\bar{B}}(\Lambda \mathbf{G}) > \bar{\kappa}(\alpha)) < 1,$$

where \mathbf{G} denotes a Gaussian random vector with mean 0 and covariance matrix I_n .

Theorem 4.3 shows that $\bar{\Phi}_{\bar{B}, \bar{\kappa}(\alpha)}$ is not subject to the zero-power trap. However, its “limiting power” $\lim_{\rho \rightarrow a} P_{\beta, \sigma, \rho}(\bar{\Phi}_{\bar{B}, \bar{\kappa}(\alpha)}) = \Pr(\bar{T}_{\bar{B}}(\Lambda \mathbf{G}) > \bar{\kappa}(\alpha))$ can in principle be low. In particular, it is always smaller than 1. This is different to the behavior of the test discussed in Section 4.1, which has limiting power equal to 1. Another limitation of Theorem 4.3 is its reliance on the additional Assumption 2.

Following up on the examples discussed in Remark 4.2, an advantage of passing from $\Phi_{B, \kappa(\alpha)}$ to $\bar{\Phi}_{\bar{B}, \bar{\kappa}(\alpha)}$, instead of passing from $\Phi_{B, \kappa(\alpha)}$ to the test discussed in Section 4.1, is that $\bar{\Phi}_{\bar{B}, \bar{\kappa}(\alpha)}$ “preserves” in some sense the optimality properties of $\Phi_{B, \kappa(\alpha)}$, but with respect to the larger group $G_{\bar{X}}$. Note, however, that this does not imply that the power functions of $\Phi_{B, \kappa(\alpha)}$ and $\bar{\Phi}_{\bar{B}, \bar{\kappa}(\alpha)}$ are “close.”

4.3. Optimality-Preserving Tests that Avoid the Zero-Power Trap

The starting point in this section is some (initial) family of tests $\varphi_\alpha : \mathbb{R}^n \rightarrow [0, 1]$ for the testing problem (3) indexed by $\alpha \in (0, 1)$. Given $\alpha \in (0, 1)$, we interpret φ_α as the (initial) test one would like to use because of some “optimality property” of its power function.² That is, the power function of φ_α

$$(\beta, \sigma, \rho) \mapsto \mathbb{E}_{\beta, \sigma, \rho}(\varphi_\alpha)$$

²What is considered optimal in a given application typically depends on the context. To ensure that the results in this section are applicable in many situations, we work with an abstract family of tests and do not restrict the “optimality property” satisfied by this family to a specific one.

is “large” for certain parameter values (β, σ, ρ) in a given subset pertaining to the alternative hypothesis $\mathbb{R}^k \times (0, \infty) \times (0, a)$.

We shall suppose that the initial test φ_α suffers from the zero-power trap, which one would like to avoid. Ideally, a test should have limiting power equal to 1, a property of the test in Section 4.1, but not of the test in Section 4.2. Furthermore, we would like to keep, at least approximately, the optimal power properties of φ_α , which was the reason why φ_α was initially considered for use. This is a property of the test in Section 4.2 (at least to some extent), but not of the test in Section 4.1. We shall now present an approach that achieves these two goals.

In what follows, we assume that the family of tests $\{\varphi_\alpha\}$ under consideration satisfies Property A, i.e., satisfies the following:

- A.1: For every $\alpha \in (0, 1)$, the test φ_α is G_X -invariant.
- A.2: For every $\alpha \in (0, 1)$, the test φ_α has size α , i.e.,

$$\sup_{\beta \in \mathbb{R}^k} \sup_{\sigma \in (0, \infty)} \mathbb{E}_{\beta, \sigma, 0}(\varphi_\alpha) = \alpha.$$

- A.3: For every $\alpha \in (0, 1)$ and every sequence $\alpha_m \in [0, \alpha]$ converging to α , we have that $\varphi_{\alpha_m}(y) \rightarrow \varphi_\alpha(y)$ holds for $\mu_{\mathbb{R}^n}$ -almost every $y \in \mathbb{R}^n$.

To illustrate the assumption, consider the following important example.

Example 4.1. Let T_B be as in (4) with B an $(n - k) \times (n - k)$ symmetric matrix such that $\lambda_1(B) < \lambda_{n-k}(B)$. For every $\alpha \in (0, 1)$, let $\kappa(\alpha)$ be the critical value from Lemma 2.3. Set φ_α equal to the nonrandomized test with rejection region $\Phi_{B, \kappa(\alpha)}$, i.e., $\varphi_\alpha := \mathbf{1}_{\Phi_{B, \kappa(\alpha)}}$. We already know that T_B is G_X -invariant and thus φ_α is G_X -invariant, for every α . Hence, A.1 is satisfied. Furthermore, from Lemma 2.3, we see that φ_α satisfies A.2. That A.3 is satisfied is an immediate consequence of the continuity of $\kappa(\cdot)$, which was established in Lemma 2.3, together with the fact that, for every $\alpha \in (0, 1)$, the set

$$\{y \in \mathbb{R}^n : T_B(y) = \kappa(\alpha)\}$$

is a $\mu_{\mathbb{R}^n}$ -null set; the latter is a consequence of Lemma B.4 in Preinerstorfer and Pötscher (2017), which shows that the cdf F , say, corresponding to $P_{0,1,0} \circ T_B$ is continuous.

Remark 4.4. While not required in Property A, typical families $\{\varphi_\alpha\}$ will also satisfy the condition that for any real numbers $\alpha_1 \leq \alpha_2$ in $(0, 1)$, it holds, for $\mu_{\mathbb{R}^n}$ -almost every $y \in \mathbb{R}^n$, that $\varphi_{\alpha_1}(y) \leq \varphi_{\alpha_2}(y)$. For instance, this is the case for the families of tests discussed in Example 4.1 (this follows from the monotonicity property of $\kappa(\cdot)$ established in Lemma 2.3). One obvious consequence of this condition is that if φ_{α_2} suffers from the zero-power trap, then φ_{α_1} suffers from the zero-power trap as well. Therefore, for such families, if φ_α suffers from the zero-power trap, there is no hope that one can easily avoid the zero-power trap by using $\varphi_{\alpha-\varepsilon}$, for some $\varepsilon > 0$ (which would at least be a test whose size does not exceed α).

Suppose in the following discussion that $k < n - 1$, that Assumption 1 holds, and that $e \notin \text{span}(X)$. Recall from Theorem 4.1 that under these conditions, the G_X -invariant test $\Phi_{C_X e e' C_X', \kappa(\alpha)}$ does not suffer from the zero-power trap, in fact has limiting power 1, at all levels $\alpha \in (0, 1)$. Using this property, we shall now define a G_X -invariant test that has approximately the same power properties of φ_α with the advantage that it has limiting power 1 just as the test $\Phi_{C_X e e' C_X', \kappa(\alpha)}$.

The basic idea is as follows (precise statements are provided further below). From Property A.3, one obtains that for $\varepsilon \in (0, \alpha)$ small, the power functions of φ_α and $\varphi_{\alpha-\varepsilon}$ are similar. Furthermore, Theorem 4.1 tells us that the test with rejection region $\Phi_{C_X e e' C_X', \kappa(\varepsilon)}$ has limiting power (as $\rho \rightarrow a$) equal to 1, and Lemma 2.3 shows that this test has size equal to ε . Hence, we could use the G_X -invariant test

$$\min(\varphi_{\alpha-\varepsilon} + \mathbf{1}_{\Phi_{C_X e e' C_X', \kappa(\varepsilon)}}, 1), \tag{12}$$

whose power function is similar to φ_α (at least for ε small), but which actually has limiting power equal to 1 (for every $0 < \varepsilon < \alpha$). Furthermore, the size of the test in equation (12) is not greater than α . However, potentially, its size is smaller than α , implying some unnecessary loss in power, which we can avoid by decreasing $\kappa(\varepsilon)$.

Instead of using the test in equation (12), we therefore suggest the G_X -invariant test

$$\varphi_{\alpha, \varepsilon}^* := \min\left(\varphi_{\alpha-\varepsilon} + \mathbf{1}_{\Phi_{C_X e e' C_X', c(\alpha, \varepsilon)}}, 1\right) = \varphi_{\alpha-\varepsilon} + (1 - \varphi_{\alpha-\varepsilon}) \mathbf{1}_{\Phi_{C_X e e' C_X', c(\alpha, \varepsilon)}}, \tag{13}$$

where $0 < c(\alpha, \varepsilon) \leq \kappa(\varepsilon)$ is chosen as the *smallest* number such that $\varphi_{\alpha, \varepsilon}^*$ has size equal to α . That such a choice of $c(\alpha, \varepsilon)$ is indeed possible is the content of the following proposition. Note that $\varphi_{\alpha, \varepsilon}^*$ is nonrandomized if the test $\varphi_{\alpha-\varepsilon}$ is nonrandomized.

PROPOSITION 4.5. *Suppose that $k < n - 1$, that $e \in \mathbb{R}^n$ satisfies $e \notin \text{span}(X)$, and that the family of tests $\{\varphi_\alpha\}$ satisfies Properties A.1 and A.2. Then, for every $\alpha \in (0, 1)$ and every $\varepsilon \in (0, \alpha)$, there exists a $c(\alpha, \varepsilon) \in (0, \kappa(\varepsilon)]$ such that*

$$\sup_{\beta \in \mathbb{R}^k} \sup_{\sigma \in (0, \infty)} \mathbb{E}_{\beta, \sigma, 0} \left[\min\left(\varphi_{\alpha-\varepsilon} + \mathbf{1}_{\Phi_{C_X e e' C_X', c(\alpha, \varepsilon)}}, 1\right) \right] = \alpha, \tag{14}$$

and such that replacing $c(\alpha, \varepsilon)$ by $c' \in (0, c(\alpha, \varepsilon))$ increases the supremum in equation (14); here, $\kappa(\varepsilon) \in (0, \|C_X e\|^2)$ denotes the unique real number such that $\Phi_{C_X e e' C_X', \kappa(\varepsilon)}$ has size equal to ε (cf. Lemma 2.3).

Note that the critical value $c(\alpha, \varepsilon)$ can be easily determined numerically by a simple line search algorithm, cf. also Remark 2.4.

Having established that the test in equation (13) is actually well-defined, we now prove that it does not suffer from the zero-power trap but has limiting power 1 for any choice of ε . Furthermore, we show that the power function of $\varphi_{\alpha, \varepsilon}^*$ approximates (even uniformly over suitable subsets of the parameter space) the

power function of φ_α as ε converges to 0. In this sense, choosing $\varepsilon > 0$ small, the test $\varphi_{\alpha,\varepsilon}^*$ at least approximately preserves the “optimal” power properties (such as, e.g., point-optimal invariance, or locally best invariance, cf. Example 4.1 above) from φ_α . Furthermore, the degree of approximation can be tuned by the user via ε .

THEOREM 4.6. *Suppose that $k < n - 1$, that Assumption 1 holds, and that $e \notin \text{span}(X)$, where e is the vector figuring in Assumption 1. Assume that the family $\{\varphi_\alpha\}$ satisfies Properties A.1 and A.2. Let $\alpha \in (0, 1)$. Then, the following holds:*

1. *For every $\varepsilon \in (0, \alpha)$, every $\beta \in \mathbb{R}^k$, and every $\sigma \in (0, \infty)$, we have*

$$\lim_{\rho \rightarrow a} \mathbb{E}_{\beta,\sigma,\rho}(\varphi_{\alpha,\varepsilon}^*) = 1;$$

in particular, $\varphi_{\alpha,\varepsilon}^$ does not suffer from the zero-power trap.*

2. *Suppose that the family $\{\varphi_\alpha\}$ also satisfies Property A.3. Let $A \subseteq [0, a)$ be such that the closure of the set*

$$\{\Sigma(\rho) / \|\Sigma(\rho)\| : \rho \in A\}$$

is contained in the set of positive definite symmetric matrices. Then,

$$\lim_{\varepsilon \rightarrow 0^+} \sup_{\beta \in \mathbb{R}^k} \sup_{\sigma \in (0, \infty)} \sup_{\rho \in A} |\mathbb{E}_{\beta,\sigma,\rho}(\varphi_{\alpha,\varepsilon}^*) - \mathbb{E}_{\beta,\sigma,\rho}(\varphi_\alpha)| = 0.$$

Remark 4.7. In the leading case, $\Sigma(\cdot)$ is a continuous function. In this case, one can choose the set A in the second part of Theorem 4.6 equal to $[0, c]$, for any $0 < c < a$ [recall that $\Sigma(\rho)$ is positive definite for every $\rho \in [0, a)$ by assumption]. Note further that since we are primarily interested in situations where the initial test φ_α suffers from the zero-power trap, while the adjusted tests $\varphi_{\alpha,\varepsilon}^*$ have limiting power 1, it is not restrictive to confine ourselves to intervals $[0, c]$ as above, as we do *not* want the power of the adjusted test to be close to the power of the initial test in a neighborhood of a . Furthermore, the optimality properties of point-optimal invariant tests (against an alternative $\bar{\rho} \in (0, a)$) or of locally best invariant tests (which are characterized by favorable power properties in the neighborhood of 0) concern only the power function over $[0, c]$ for a suitably chosen $c < a$.

Remark 4.8. The tuning parameter ε needs to be chosen by the user in each particular application. In principle, the user can plot the power functions for various values of ε and can then decide upon inspection, which value of ε provides the best solution. For a specific example, we refer to Section 5 below.

Remark 4.9. Finally, we point out that the construction of $\varphi_{\alpha,\varepsilon}^*$ in equation (13) and the conditions in Proposition 4.5 and Theorem 4.6 do not require the initial test φ_α to suffer from the zero-power trap. While this is clearly our main focus, this observation shows that our method can also be applied in case the limiting-power of φ_α is greater than 0 but smaller than 1. In such a situation, using $\varphi_{\alpha,\varepsilon}^*$ instead of φ_α can be advantageous as well.

5. NUMERICAL RESULTS

In order to illustrate and compare the power properties of the tests introduced in Section 4, we now consider a simple example from spatial econometrics, in which the zero-power trap occurs for a popular test (cf. also the discussion of Kelly (2019) in the introduction). We focus on a situation where the correlation between the observations is a consequence of their proximity, which might be spatial, but could also be, e.g., social, and which is encoded in the adjacency (“weights”) matrix of a graph.

One important model in this case is the spatial (autoregressive) error model, which leads to

$$\Sigma(\rho) = [(I - \rho W')(I - \rho W)]^{-1},$$

for W a fixed weights matrix, which is assumed to be (elementwise) nonnegative and irreducible with zeros on the main diagonal. By the Perron–Frobenius theorem (e.g., Horn and Johnson (1985), Theorem 8.4.4), the matrix W then has a positive (real) eigenvalue $\lambda_{\max}(W)$, say, with algebraic multiplicity (and thus also geometric multiplicity) equal to 1, such that any other real or complex zero of the characteristic polynomial of W is in absolute value not larger than $\lambda_{\max}(W)$. We assume that the parameter $\rho \in [0, \lambda_{\max}^{-1}(W))$, i.e., $a = \lambda_{\max}^{-1}(W)$. For f_{\max} a normalized eigenvector of W w.r.t. $\lambda_{\max}(W)$, it is not too difficult to see that Assumption 1 is satisfied (with $e = f_{\max}$), and that Assumption 2 is satisfied. For details, we refer to Section 4.1 in Preinerstorfer and Pötscher (2017).

The model depends, besides the design matrix X , on the specific form of the weights matrix W , which encodes the dependence relation of the observations. Subsequently, we reconsider a simple example considered in Section 3 of Krämer (2005), who has observed (cf. his Figure 1) that for a weights matrix derived by the Queen criterion from a 4×4 regular lattice, and for $X = (1, \dots, 1)' \in \mathbb{R}^{16}$, the Cliff–Ord test suffers from the zero-power trap for $\alpha = 5\%$. We recall that the Cliff–Ord test is based on a test statistic as in equation (4) and with $B = C_X(W + W')C_X'$.

The power function of the Cliff–Ord test and the power functions of the tests described in Section 4 were obtained numerically (cf. also Remark 2.4), and are shown in Figure 1. The figure also shows the power *envelope* in the class of G_X -invariant tests. That is, for each alternative $\bar{\rho} \in (0, \lambda_{\max}(W)^{-1})$, Figure 1 shows the power of the point-optimal G_X -invariant level $\alpha = 5\%$ test against the alternative $\bar{\rho}$. Recall from Remark 2.2 that the point-optimal invariant test against alternative $\bar{\rho}$ is based on a test statistic as in (4) and with $B = -[C_X \Sigma(\bar{\rho}) C_X']^{-1}$. In this example, the power envelope is not attained by any G_X -invariant test, but it serves the purpose of providing an upper bound for comparison.

While Figure 1 illustrates that the approaches discussed in Sections 4.1 and 4.2 avoid the zero-power trap, it reveals at the same time that the power functions of these tests are not completely satisfying. On the one hand, even though the test introduced in Section 4.1 does not suffer from the zero-power trap, it has low power in a large portion of the alternative. On the other hand, the test from Section 4.2




FIGURE 1. Comparison of power functions. The horizontal line corresponds to $\alpha = .05$. “Env.” corresponds to the power envelope; “CO test” to the power function of the Cliff–Ord test; “Sol. 1” corresponds to the power function of the test from Section 4.1; “CO Sol. 2” corresponds to the power function of the artificial regressor-based Cliff–Ord test as discussed in Section 4.2; “eps = .01” corresponds to the test in Section 4.3 with φ_α the size α Cliff–Ord test and $\varepsilon = .01$; the remaining “eps = .006” and “eps = .002” correspond likewise to the tests in Section 4.3, but for the corresponding values of ε .

based on the Cliff–Ord test (i.e., as in equation (11) with $\bar{B} = C_{(X,e)}(W + W')C'_{(X,e)}$) with artificial regressor $e = f_{\max}$ avoids the zero-power trap as well, and has a power function that practically coincides with the power envelope for small values of ρ . But its limiting power is smaller than 1 (in fact is only 0.619).

Figure 1 also contains the power function of some tests corresponding to the procedure outlined in Section 4.3 applied to the family φ_α of level- α Cliff–Ord tests (cf. Example 4.1). It shows the power functions corresponding to $\varepsilon \in \{.002, .006, .01\}$. These three tests have very good power properties. For small values of ρ , their power functions are practically identical to the power function of the Cliff–Ord test, and hence to the power envelope. But for larger values of ρ , their power functions are much closer to the power envelope than that of the Cliff–Ord test. In particular, by construction, their power converges to 1 as ρ gets close to a . One can also observe that smaller values of ε lead to power functions that are closer to the power function of the Cliff–Ord test for ρ close to 0, whereas larger values of ε lead to power functions that are closer to the power envelope for ρ close to a .

6. CONCLUSION

In the present article, we have reconsidered the zero-power trap phenomenon in testing for correlation in a general framework. Most importantly, we have suggested a way to construct “approximately optimal tests” that avoid the trap. For practical purposes, if an initial test, such as the Cliff–Ord test in the example discussed in Section 5, turns out to suffer from the zero-power trap, we suggest to use the method introduced in Section 4.3 to obtain a modified test with the following properties: (i) it has a similar power function as the initial test, (ii) it does not suffer from the zero-power trap, and (iii) its limiting power equals 1. The tuning parameter ε involved in the construction of the modified test can be chosen by graphically comparing the power functions of modified tests corresponding to different values of the tuning parameter with the power envelope and the power function of the initial test. The heuristic underlying our construction can be interpreted as a finite sample variant of the power enhancement principle of Fan et al. (2015); more specifically, our construction is based on the formulation of the power enhancement principle developed in Section 3 of Kock and Preinerstorfer (2019). The approach, which is not restricted to the testing problem under consideration, might be of some interest in its own right.

APPENDICES

A. Proofs for Results in Section 1

Proof of Lemma 2.3. Lemma B.4 in Preinerstorfer and Pötscher (2017) shows that the cdf F , say, corresponding to $P_{0,1,0} \circ T_B$ is continuous, that $F(\lambda_1(B)) = 0$, $F(\lambda_{n-k}(B)) = 1$, and that F is strictly increasing on $[\lambda_1(B), \lambda_{n-k}(B)]$. Hence, the function $f : [\lambda_1(B), \lambda_{n-k}(B)] \rightarrow [0, 1]$ defined via

$$c \mapsto P_{0,1,0}(\Phi_{B,c}) = 1 - F(c)$$

is continuous, strictly decreasing, and satisfies $f(\lambda_1(B)) = 1$ and $f(\lambda_{n-k}(B)) = 0$. Set $\kappa = f^{-1}$, i.e., the inverse of f , which is continuous, strictly decreasing, and obviously satisfies $\kappa(0) = \lambda_{n-k}(B)$ and $\kappa(1) = \lambda_1(B)$. Then $P_{0,1,0}(\Phi_{B,\kappa(\alpha)}) = \alpha$, for every $\alpha \in [0, 1]$. Finally, recall that T_B is G_X -invariant, from which it follows (cf. Remark 2.3 in Preinerstorfer and Pötscher (2017)) that, for every $c \in \mathbb{R}$, every $\beta \in \mathbb{R}^k$, and every $\sigma \in (0, \infty)$, we have $P_{\beta,\sigma,0}(\Phi_{B,c}) = P_{0,1,0}(\Phi_{B,c})$. Hence, $P_{\beta,\sigma,0}(\Phi_{B,\kappa(\alpha)}) = \alpha$ holds for every $\beta \in \mathbb{R}^k$, every $\sigma \in (0, \infty)$, and every $\alpha \in [0, 1]$. The uniqueness part is obvious. \square

B. Proofs for Results in Section 3

Proof of Theorem 3.1. We apply Theorem 2.7 in Preinerstorfer and Pötscher (2017). Their Assumption 1 coincides with ours and is thus satisfied. Furthermore, by our Gaussianity assumption, their Assumption 3 is satisfied in our framework (with \mathbf{z} a normally distributed random vector with mean 0 and covariance matrix I_n), and we can use Part 1 of their Proposition 2.6 to conclude that their Assumption 2 is satisfied. The statement now follows from Theorem 2.7 in Preinerstorfer and Pötscher (2017) for

the special case $\varphi(e) = 0$. The last statement follows from Remark 2.8(i) in the same reference. \square

Proof of Theorem 3.2. We use Corollary 2.21 in Preinerstorfer and Pötscher (2017). That their Assumptions 1 and 2 are satisfied follows as in the proof of Theorem 3.1 above. Recall from Lemma 2.3 that κ is a strictly decreasing and continuous bijection from $[0, 1]$ to $[\lambda_1(B), \lambda_{n-k}(B)]$, implying that, for $\alpha \in (0, 1)$, we have $\kappa(\alpha) \in (\lambda_1(B), \lambda_{n-k}(B))$. We can hence apply Corollary 2.21 in Preinerstorfer and Pötscher (2017) to conclude that (under our assumptions), for $\alpha \in (0, 1)$ such that $T_B(e) < \kappa(\alpha)$, we have (7). \square

Proof of Lemma 3.3. Noting that both $e \notin \text{span}(X)$ and $\lambda_1(B(X)) < \lambda_{n-k}(B(X))$ follow from $C_X e \notin \text{Eig}(B(X), \lambda_{n-k}(B(X)))$, Condition (8) together with the definition of $T_{B(X)}$ in equation (4) can be used to verify $\lambda_1(B(X)) \leq T_{B(X)}(e) < \lambda_{n-k}(B(X))$. Thus, Lemma 2.3 gives $P_{0,1,0}^X(\Phi_{B(X), T_{B(X)}(e)}) \in (0, 1]$ and $T_{B(X)}(e) < \kappa(\alpha)$ for every $\alpha \in (0, P_{0,1,0}^X(\Phi_{B(X), T_{B(X)}(e)}))$. We can now apply Theorem 3.2 to conclude. \square

LEMMA B.1. Let $M \in \mathbb{R}^{n \times n}$ be symmetric, let $v \in \mathbb{R}^n$ be such that $\|v\| = 1$, and suppose that $1 \leq d < n - 1$. Then,

$$\mathcal{D}(n, d) := \{L \in \mathbb{R}^{n \times d} : \text{rank}(L) = d,$$

$$\Pi_{\text{span}(L)^\perp} v \text{ is an eigenvector of } \Pi_{\text{span}(L)^\perp} M \Pi_{\text{span}(L)^\perp}\}$$

can be written as

$$\{L \in \mathbb{R}^{n \times d} : \det(L'L) \neq 0, \|\Pi_{\text{span}(L)^\perp} v\| \neq 0\} \cap \{L \in \mathbb{R}^{n \times d} : p_M(L) = 0\}, \tag{B.1}$$

for $p_M : \mathbb{R}^{n \times d} \rightarrow \mathbb{R}$ a multivariate polynomial, which is given in the proof. Furthermore, $p_M \equiv 0$ if and only if $M = c_1 I_n + c_2 v v'$ holds for real numbers c_1 and c_2 .

Proof. We first establish that $\mathcal{D}(n, d)$ equals the set in equation (B.1). To this end, let $L \in \mathbb{R}^{n \times d}$ satisfy $\text{rank}(L) = d$, or equivalently $\det(L'L) \neq 0$. If $\Pi_{\text{span}(L)^\perp} v = 0$, the vector $\Pi_{\text{span}(L)^\perp} v$ cannot be an eigenvector of $\Pi_{\text{span}(L)^\perp} M \Pi_{\text{span}(L)^\perp}$. If $\Pi_{\text{span}(L)^\perp} v \neq 0$, $\Pi_{\text{span}(L)^\perp} v$ is an eigenvector of the symmetric matrix $\Pi_{\text{span}(L)^\perp} M \Pi_{\text{span}(L)^\perp}$ if and only if

$$\text{rank}\left(\left(\Pi_{\text{span}(L)^\perp} v, \Pi_{\text{span}(L)^\perp} M \Pi_{\text{span}(L)^\perp} v\right)\right) < 2. \tag{B.2}$$

We can write this rank condition equivalently as

$$0 = \det\left[\left(\Pi_{\text{span}(L)^\perp} v, \Pi_{\text{span}(L)^\perp} M \Pi_{\text{span}(L)^\perp} v\right)' \left(\Pi_{\text{span}(L)^\perp} v, \Pi_{\text{span}(L)^\perp} M \Pi_{\text{span}(L)^\perp} v\right)\right]. \tag{B.3}$$

Writing $\Pi_{\text{span}(L)^\perp} = I_n - \det(L'L)^{-1} L \text{adj}(L'L) L'$ (throughout we use the convention that the adjoint of a 1×1 matrix equals 1), and premultiplying (B.3) by $\det(L'L)^8 \neq 0$, one sees that (B.3) is equivalent to

$$0 = \det\left[\left(\det(L'L) Q(L)v, Q(L) M Q(L)v\right)' \left(\det(L'L) Q(L)v, Q(L) M Q(L)v\right)\right] =: p_M(L),$$

where $Q(L) := \det(L'L) I_n - L \text{adj}(L'L) L'$. Note that $L \mapsto p(L)$ defines a multivariate polynomial on $\mathbb{R}^{n \times d}$. It follows that $\mathcal{D}(n, d)$ has the claimed form.

We now prove that $p_M \equiv 0$ if and only if $M = c_1 I_n + c_2 v v'$ holds for real numbers c_1 and c_2 . To this end, note that if M is of the specific form $c_1 I_n + c_2 v v'$ for real numbers c_1 and c_2 , one has, for every $L \in \mathbb{R}^{n \times d}$, that

$$\Pi_{\text{span}(L)^\perp} M \Pi_{\text{span}(L)^\perp} v = (c_1 + c_2 v' \Pi_{\text{span}(L)^\perp} v) \Pi_{\text{span}(L)^\perp} v. \tag{B.4}$$

For L such that $\det(L'L) \neq 0$, the statement $p_M(L) = 0$ is equivalent to (B.2), which follows from (B.4). If L satisfies $\det(L'L) = 0$, we obviously have $p_M(L) = 0$. Thus, $p_M \equiv 0$, for all M of this specific form.

Now assume that M cannot be written as $c_1 I_n + c_2 v v'$ for real numbers c_1 and c_2 . To show that $p_M \neq 0$, it suffices to construct a single L such that $p_M(L) \neq 0$. We consider two cases (a) and (b):

(a) We first show that one can find an L such that $p_M(L) \neq 0$ in the special case where v is not an eigenvector of M . Let u_1, \dots, u_n be an orthonormal basis of eigenvectors of M with corresponding eigenvalues $\lambda_1(M), \dots, \lambda_n(M)$. Note that there then exist two indices $j \neq l$, say, such that $\lambda_j(M) \neq \lambda_l(M)$ and such that $v'u_j \neq 0$ and $v'u_l \neq 0$ (otherwise, v would be an eigenvector of M ; recall that $v \neq 0$).

Now, define the matrix $L_\perp = (u_j, u_l, z_1, \dots, z_{n-d-2})$ for z_1, \dots, z_{n-d-2} linearly independent elements of $\text{span}(u_j, u_l, v)^\perp$ (with the convention that $L_\perp = (u_j, u_l)$ if $n - d = 2$; note that $n - d \geq 2$ holds by assumption). Such a choice of z_1, \dots, z_{n-d-2} is possible as $d \geq 1$ by assumption. Note that $\text{rank}(L_\perp) = n - d$.

Next, let L be an $n \times d$ matrix with $\text{span}(L) = \text{span}(L_\perp)^\perp$. Then, L is of full column rank, and $\Pi_{\text{span}(L)^\perp} v \neq 0$. From the discussion preceding the definition of p_M , we see that it thus remains to verify that $\Pi_{\text{span}(L)^\perp} v$ is not an eigenvector of $\Pi_{\text{span}(L)^\perp} M \Pi_{\text{span}(L)^\perp}$. But $\Pi_{\text{span}(L)^\perp} v = \Pi_{\text{span}(u_j, u_l)} v = u'_j v u_j + u'_l v u_l$, implying $\Pi_{\text{span}(L)^\perp} M \Pi_{\text{span}(L)^\perp} v = \lambda_j(M) u'_j v u_j + \lambda_l(M) u'_l v u_l$. Hence, if $\Pi_{\text{span}(L)^\perp} v$ was an eigenvector of $\Pi_{\text{span}(L)^\perp} M \Pi_{\text{span}(L)^\perp}$, we would have

$$\lambda_j(M) u'_j v u_j + \lambda_l(M) u'_l v u_l = c (u'_j v u_j + u'_l v u_l),$$

for some $c \in \mathbb{R}$, which gives the contradiction $\lambda_j(M) = \lambda_l(M) = c$.

(b) Next, we consider the case where v is an eigenvector of M to the eigenvalue $\lambda_i(M)$, say. Let u_1, \dots, u_n be an orthonormal basis of eigenvectors of M corresponding to its eigenvalues $\lambda_1(M), \dots, \lambda_n(M)$, and where $u_i = v$ holds. By assumption, M is not of the form $c_1 I_n + c_2 v v'$. Together with v being an eigenvector of M , this implies (via a diagonalization argument) existence of two indices j and l , say, such that i, j, l are pairwise distinct and such that $\lambda_j(M) \neq \lambda_l(M)$.

Now, define $L_\perp = (x, y, z_1, \dots, z_{n-d-2})$, where $x = v + u_j$ and $y = v + u_l$, and where z_1, \dots, z_{n-d-2} are linearly independent elements of $\text{span}(u_j, u_l, v)^\perp$ (with the convention that $L_\perp = (x, y)$ if $n - d = 2$; recall that $n - d \geq 2$ holds by assumption). Such a construction is possible as $d \geq 1$ by assumption. Note that $\text{rank}(L_\perp) = n - d$.

Define L as an $n \times d$ matrix with $\text{span}(L) = \text{span}(L_\perp)^\perp$. Then, L is of full column rank, and $\Pi_{\text{span}(L)^\perp} v \neq 0$. Arguing as in (a), it now remains to verify that $\Pi_{\text{span}(L)^\perp} v$ is not an eigenvector of $\Pi_{\text{span}(L)^\perp} M \Pi_{\text{span}(L)^\perp}$: It is easy to see that

$$\Pi_{\text{span}(L)^\perp} v = \Pi_{\text{span}(x, y)} v = 3^{-1} (x + y),$$

and that, using the expression in the previous display and a simple computation,

$$\begin{aligned} & \Pi_{\text{span}(L)^\perp} M \Pi_{\text{span}(L)^\perp} v \\ &= 9^{-1} [(2\lambda_i(M) + 2\lambda_j(M) - \lambda_l(M))x + (2\lambda_i(M) - \lambda_j(M) + 2\lambda_l(M))y]. \end{aligned}$$

Hence, for this choice of L , the vector $\Pi_{\text{span}(L)^\perp} v$ is an eigenvector of $\Pi_{\text{span}(L)^\perp} M \Pi_{\text{span}(L)^\perp}$ if and only if

$$3^{-1}(x + y) = c 9^{-1} [(2\lambda_i(M) + 2\lambda_j(M) - \lambda_l(M))x + (2\lambda_i(M) + 2\lambda_l(M) - \lambda_j(M))y], \tag{B.5}$$

for some $c \in \mathbb{R}$. The number c must then necessarily be nonzero. But this implies (premultiply both sides of (B.5) by u'_j , then by u'_i , and compare the two equations obtained) that $\lambda_j(M) = \lambda_l(M)$, a contradiction. □

Proof of Proposition 3.4. We start with the claim that up to a subset of a $\mu_{\mathbb{R}^{n \times k}}$ -null set of exceptional matrices, every $X \in \mathbb{R}^{n \times k}$ satisfies (8). From $k < n$, it follows that $\mu_{\mathbb{R}^{n \times k}}(\{X \in \mathbb{R}^{n \times k} : \text{rank}(X) < k\}) = 0$. Hence, it suffices to show that

$$\{X \in \mathbb{R}^{n \times k} : \text{rank}(X) = k \text{ and } C_X e \in \text{Eig}(B(X), \lambda_{n-k}(B(X)))\} \tag{B.6}$$

is contained in a $\mu_{\mathbb{R}^{n \times k}}$ -null set. To show this, we separately consider two cases (a) and (b):

(a) Suppose first that $M = c_1 I_n + c_2 e e'$ for real numbers c_1, c_2 where $c_2 < 0$. Then, we shall show that the set in equation (B.6) simplifies to

$$\{X \in \mathbb{R}^{n \times k} : \text{rank}(X) = k \text{ and } e \in \text{span}(X)\}. \tag{B.7}$$

To see this, note that in this case and for $X \in \mathbb{R}^{n \times k}$ so that $\text{rank}(X) = k$, we have

$$\text{Eig}(B(X), \lambda_{n-k}(B(X))) = \text{Eig}(C_X M C'_X, \lambda_{n-k}(C_X M C'_X)) = \text{span}(C_X e)^\perp,$$

where we used the assumption in (9) to obtain the first equality, and the specific structure of M and $c_2 < 0$ to obtain the second equality. Thus, $C_X e \in \text{Eig}(B(X), \lambda_{n-k}(B(X)))$ is possible only if $C_X e = 0$, which is equivalent to $e \in \text{span}(X)$. Therefore, (B.6) simplifies to (B.7). But, by assumption, $1 \leq k < n$ holds, from which it is easy to see, noting that $\|e\| = 1$, that $\mu_{\mathbb{R}^{n \times k}}(\{X \in \mathbb{R}^{n \times k} : e \in \text{span}(X)\}) = 0$. Therefore, the set in (B.7), and equivalently the set in (B.6), is a $\mu_{\mathbb{R}^{n \times k}}$ -null set in this case.

(b) Consider now the case where M is not a linear combination of I_n and $e e'$. Using equation (9), we can write the set defined in (B.6) equivalently as

$$\{X \in \mathbb{R}^{n \times k} : \text{rank}(X) = k \text{ and } C_X e \in \text{Eig}(C_X M C'_X, \lambda_{n-k}(C_X M C'_X))\}. \tag{B.8}$$

For $X \in \mathbb{R}^{n \times k}$ of full column rank, the property $C'_X C_X = \Pi_{\text{span}(X)^\perp} = \Pi_{\text{span}(X)^\perp}^2$ can be used to verify that

$$C_X e \in \text{Eig}(C_X M C'_X, \lambda_{n-k}(C_X M C'_X))$$

implies

$$\Pi_{\text{span}(X)^\perp} e \in \text{Eig}(\Pi_{\text{span}(X)^\perp} M \Pi_{\text{span}(X)^\perp}, \lambda_{n-k}(C_X M C'_X)).$$

Thus, if $e \notin \text{span}(X)$, then $\Pi_{\text{span}(X)^\perp} e \neq 0$, and $C_X e \in \text{Eig}(C_X M C'_X, \lambda_{n-k}(C_X M C'_X))$ implies that $\Pi_{\text{span}(X)^\perp} e$ is an eigenvector of $\Pi_{\text{span}(X)^\perp} M \Pi_{\text{span}(X)^\perp}$. Thus, the set in equation (B.8)

is contained in the union of the $\mu_{\mathbb{R}^{n \times k}}$ -null set $\{X \in \mathbb{R}^{n \times k} : e \in \text{span}(X)\}$ and the set $\{X \in \mathbb{R}^{n \times k} : \text{rank}(X) = k, \Pi_{\text{span}(X)^\perp} e \text{ is an eigenvector of } \Pi_{\text{span}(X)^\perp} M \Pi_{\text{span}(X)^\perp}\}$.

(B.9)

It thus suffices to verify that the set in (B.9) is a $\mu_{\mathbb{R}^{n \times k}}$ -null set. Lemma B.1 (applied with $k = d$ and $v = e$) shows that (B.9) is the subset of an algebraic set. Note that the assumptions in Lemma B.1 are satisfied as $1 \leq k < n - 1$ is assumed. The lemma also provides the information that a multivariate polynomial defining this algebraic set does not vanish everywhere. Hence, it follows that the set in (B.9) is contained in a $\mu_{\mathbb{R}^{n \times k}}$ -null set. This finishes Case (b).

We now prove the two remaining claims concerning $\mathcal{X}(\alpha; B)$. For the monotonicity claim: If $\mathcal{X}(\alpha_2; B)$ is empty, there is nothing to prove. Consider the case where $\mathcal{X}(\alpha_2; B) \neq \emptyset$. Let $X \in \mathcal{X}(\alpha_2; B)$. By definition of $\mathcal{X}(\alpha_2; B)$, the matrix X has full column rank and $\lambda_1(B(X)) < \lambda_{n-k}(B(X))$. From $0 < \alpha_1 \leq \alpha_2 < 1$, it thus follows from Lemma 2.3 that $\kappa(\alpha_2) \leq \kappa(\alpha_1)$. Hence, $\Phi_{B(X), C_X, \kappa(\alpha_1)} \subseteq \Phi_{B(X), C_X, \kappa(\alpha_2)}$ and one obtains $X \in \mathcal{X}(\alpha_1; B)$. Finally, note that Lemma 3.3 shows that if X satisfies (8), then $X \in \bigcup_{m \in \mathbb{N}} \mathcal{X}(\alpha_m; B)$. The first (already established) part of the current proposition hence proves the last claim. \square

LEMMA B.2. Let $M \in \mathbb{R}^{n \times n}$ be symmetric, let $v \in \mathbb{R}^n$ be such that $\|v\| = 1$, and suppose that M cannot be written as $c_1 I_n + c_2 v v'$ for real numbers c_1, c_2 where $c_2 \geq 0$. Let $d \in \mathbb{N}$ be such that $d < n - 1$. Then:

- 1. There exist a sequence $L_m \in \mathbb{R}^{n \times d}$, such that $L'_m L_m = I_d$ and $L_m \rightarrow L^*$ as $m \rightarrow \infty$, a vector $u \in \mathbb{R}^n$ with $\|u\| = 1$, and a real number $c > \lambda_{\min}(M)$, such that: $\Pi_{\text{span}(L_m)^\perp} v \neq 0$ and $\Pi_{\text{span}(L_m)^\perp} u \neq 0$ holds, for every $m \in \mathbb{N}$, such that

$$\lim_{m \rightarrow \infty} v' \Pi_{\text{span}(L_m)^\perp} M \Pi_{\text{span}(L_m)^\perp} v / v' \Pi_{\text{span}(L_m)^\perp} v = \lambda_{\min}(M), \tag{B.10}$$

and such that, for every $m \in \mathbb{N}$, we have

$$u' \Pi_{\text{span}(L_m)^\perp} M \Pi_{\text{span}(L_m)^\perp} u / u' \Pi_{\text{span}(L_m)^\perp} u = c. \tag{B.11}$$

- 2. Let B be a function from the set of full column rank $n \times d$ matrices to the set of symmetric $(n - d) \times (n - d)$ -dimensional matrices. Suppose there exists a function F from the set of $(n - d) \times (n - d)$ matrices to itself, such that, for every $L \in \mathbb{R}^{n \times d}$ of full column rank, $B(L) = F(C_L M C'_L)$ holds for a suitable choice of $C_L \in \mathbb{R}^{(n-d) \times n}$ satisfying $C_L C'_L = I_{n-d}$ and $C'_L C_L = \Pi_{\text{span}(L)^\perp}$. Suppose further that F is continuous at every element A , say, of the closure of $\{C_L M C'_L : L \in \mathbb{R}^{n \times d}, \text{rank}(L) = d\} \subseteq \mathbb{R}^{(n-d) \times (n-d)}$, and that, for every such A , we have

$$\text{Eig}(F(A), \lambda_1(F(A))) = \text{Eig}(A, \lambda_1(A)). \tag{B.12}$$

Then, the sequence L_m obtained in Part 1 satisfies $C_{L_m} v \neq 0$, for every $m \in \mathbb{N}$,

$$\lim_{m \rightarrow \infty} \left[v' C'_{L_m} B(L_m) C_{L_m} v / \|C_{L_m} v\|^2 - \lambda_1(B(L_m)) \right] = 0, \tag{B.13}$$

and

$$\liminf_{m \rightarrow \infty} \left[\lambda_{n-k}(B(L_m)) - \lambda_1(B(L_m)) \right] = \delta, \tag{B.14}$$

for some positive real number δ .

Proof. Before we prove Part 1, we note that it suffices to verify the existence claim without the requirement that L_m converges: Convergence of L_m can then be achieved by passing to a subsequence. The proof of Part 1 is divided in the two cases (1.a) and (1.b).

(1.a) Consider first the case where $v \in \text{Eig}(M, \lambda_{\min}(M))$: Let $u \in \text{Eig}(M, \lambda_{\max}(M))$ such that $\|u\| = 1$, and set $L_{m,\perp} := (u, v, w_1, \dots, w_{n-d-2})$ for w_1, \dots, w_{n-d-2} linearly independent elements of $\text{span}((u, v))^\perp$ (with the implicit understanding that $L_{m,\perp} = (u, v)$ in case $d = n - 2$). By assumption M is not a multiple of I_n , thus $\lambda_{\min}(M) < \lambda_{\max}(M)$, from which it also follows that $L_{m,\perp}$ has full column rank $n - d \geq 2$, for every $m \in \mathbb{N}$. For every $m \in \mathbb{N}$, set L_m equal to an $n \times d$ matrix such that $L'_m L_m = I_d$ and $\text{span}(L_m)^\perp = \text{span}(L_{m,\perp})$. Then, equations (B.10) and (B.11) (with $c = \lambda_{\max}(M)$) follow immediately from $\Pi_{\text{span}(L_m)^\perp} v = v$ and $\Pi_{\text{span}(L_m)^\perp} u = u$.

(1.b) Next, we consider the case where $v \notin \text{Eig}(M, \lambda_{\min}(M))$: We first claim that there must exist an $x \in \text{Eig}(M, \lambda_{\min}(M))$ such that $\|x\| = 1$ and a vector $u \in \text{span}(v, x)^\perp$ such that $\|u\| = 1$ and such that $u' M u > \lambda_{\min}(M)$. We argue by contradiction: First of all, if the claim was false, then $\dim(\text{Eig}(M, \lambda_{\min}(M))) = n - 1$ would follow. We could then choose v_1, \dots, v_{n-1} an orthonormal basis of $\text{Eig}(M, \lambda_{\min}(M))$. Under the assumption that the above claim was wrong, it would further follow that $\text{span}(v, v_i)^\perp \subseteq \text{Eig}(M, \lambda_{\min}(M))$, for every $i = 1, \dots, n - 1$, implying $\text{span}(v, v_i)^\perp \subseteq \text{span}(v_1, \dots, v_{i-1}, v_{i+1}, \dots, v_{n-1})$, for every $i = 1, \dots, n - 1$, which, by a dimension argument using $v \notin \text{Eig}(M, \lambda_{\min}(M))$, is equivalent to

$$\text{span}(v, v_i)^\perp = \text{span}(v_1, \dots, v_{i-1}, v_{i+1}, \dots, v_{n-1}) \quad \text{for } i = 1, \dots, n - 1;$$

or equivalently

$$\text{span}(v, v_i) = \text{span}(v_1, \dots, v_{i-1}, v_{i+1}, \dots, v_{n-1})^\perp \quad \text{for } i = 1, \dots, n - 1.$$

Since $n \geq 3$, setting $i = 1$ and $i = 2$ in the previous display then shows that v is orthogonal to v_1, \dots, v_{n-1} , and hence $\text{span}(v) = \text{Eig}(M, \lambda_{\max}(M))$ would follow. But then, we could conclude that $M = \lambda_{\min}(M)I_n + (\lambda_{\max}(M) - \lambda_{\min}(M))vv'$, a contradiction. Now, let $x \in \text{Eig}(M, \lambda_{\min}(M))$ be such that $\|x\| = 1$, and choose a corresponding $u \in \text{span}(v, x)^\perp$ such that $\|u\| = 1$ and such that $u' M u > \lambda_{\min}(M)$. Let $b_m \neq 0$ be a sequence that converges to 0 and such that $b_m \neq -v'x$ holds, for every $m \in \mathbb{N}$. Then, we define $v_m := x + b_m v \perp u$ and set $L_{m,\perp} := (u, v_m, w_1, \dots, w_{n-d-2})$ (with $L_{m,\perp} = (u, v_m)$ in case $d = n - 2$), for w_1, \dots, w_{n-d-2} linearly independent elements of $\text{span}(u, v, x)^\perp$ (which is possible as $d \geq 1$). As $v_m \neq 0$ follows from $b_m \neq -v'x$, the matrix $L_{m,\perp}$ has full column rank $n - d \geq 2$, for every $m \in \mathbb{N}$. Now, for every $m \in \mathbb{N}$, set L_m equal to an $n \times d$ matrix such that $L'_m L_m = I_d$ and $\text{span}(L_m)^\perp = \text{span}(L_{m,\perp})$. Then,

$$\Pi_{\text{span}(L_m)^\perp} v = \Pi_{\text{span}(L_{m,\perp})} v = \Pi_{\text{span}((u, v_m))} v = \Pi_{\text{span}((v_m))} v = a_m v_m,$$

where $a_m = (v'x + b_m)/(v'_m v_m) \neq 0$ holds, for all m . From $v_m \neq 0$, we thus obtain $\Pi_{\text{span}(L_m)^\perp} v \neq 0$, for every $m \in \mathbb{N}$. But $v_m \rightarrow x$ (as $m \rightarrow \infty$) hence shows that

$$a_m^{-2} v' \Pi_{\text{span}(L_m)^\perp} M \Pi_{\text{span}(L_m)^\perp} v \rightarrow \lambda_{\min}(M) \quad \text{and} \quad a_m^{-2} v' \Pi_{\text{span}(L_m)^\perp} v \rightarrow 1,$$

which implies (B.10). Equation (B.11) now follows (with $c = u' M u$), because $u \in \text{span}(L_m)^\perp$ gives $\Pi_{\text{span}(L_m)^\perp} u = u$, and since u was chosen such that $\|u\| = 1$ and $u' M u > \lambda_{\min}(M)$.

This establishes Part 1. We now prove Part 2.

(2) Obviously, $C_{L_m} v \neq 0$ follows from $\Pi_{\text{span}(L_m)^\perp} v \neq 0$. Consider first equation (B.13). Let m' be an arbitrary subsequence of m . Define $v_m := C_{L_m} v / \|C_{L_m} v\|$ and $A_m := C_{L_m} M C_{L_m}'$. Clearly, $\|v_m\| = 1$, and A_m is a norm-bounded sequence, because $C_{L_m} C_{L_m}' = I_{n-d}$. The latter also implies

$$\lambda_1(M) \leq \lambda_1(A_m) \leq \lambda_{n-d}(A_m) \leq \lambda_n(M) \quad \text{for every } m \in \mathbb{N}. \tag{B.15}$$

Hence, we can choose a subsequence m'' of m' , say, along which v_m and A_m converge to v_* and A , say, respectively. Note that $\|v_*\| = 1$. Next, we use $C_{L_m}' C_{L_m} = \Pi_{\text{span}(L_m)^\perp}$ to rewrite

$$v' \Pi_{\text{span}(L_m)^\perp} M \Pi_{\text{span}(L_m)^\perp} v / v' \Pi_{\text{span}(L_m)^\perp} v = v'_m C_{L_m}' M C_{L_m}' v_m = v'_m A_m v_m,$$

and use equation (B.10) to conclude that along m'' we have $v'_m A_m v_m \rightarrow v'_* A v_* = \lambda_{\min}(M)$. From equation (B.15), we obtain $\lambda_{\min}(M) = \lambda_{\min}(A)$, hence

$$v_* \in \text{Eig}(A, \lambda_1(A)) = \text{Eig}(F(A), \lambda_1(F(A))),$$

where the equality is obtained from (B.12). Finally, we observe that along m'' we have (using continuity of F) that $B(L_m) = F(A_m) \rightarrow F(A)$, from which

$$v' C_{L_m}' B(L_m) C_{L_m} v / \|C_{L_m} v\|^2 = v'_m F(A_m) v_m \rightarrow v'_* F(A) v_* = \lambda_1(F(A)),$$

and $\lambda_1(B(L_m)) \rightarrow \lambda_1(F(A))$ follows (along m''). Hence, we have shown that the statement in equation (B.13) holds along the subsequence m'' of m' . But m' was arbitrary. Therefore, we are done.

For (B.14), we argue by contradiction. Note first that the limit inferior in (B.14) cannot be infinite, because $B(L_m) = F(C_{L_m} M C_{L_m}')$, and because of the continuity property of F together with boundedness of $C_{L_m} M C_{L_m}'$. Now, assuming (B.14) were false, we could choose a subsequence m' of m such that $\lambda_{n-k}(B(L_{m'})) - \lambda_1(B(L_{m'})) \rightarrow 0$. Choose a subsequence m'' of m' along which v_m just defined above, $u_m := C_{L_m} u / \|C_{L_m} u\|$ (note that $C_{L_m} u \neq 0$ follows from $\Pi_{\text{span}(L_m)^\perp} u \neq 0$), and $A_m := C_{L_m} M C_{L_m}'$ converge to v_* , u_* , and A , respectively (where v_* and A might differ from the limits in the preceding paragraph where we established equation (B.13)). Note also that $\|v_*\| = \|u_*\| = 1$. Recall that $B(L_m) = F(A_m)$, and note that

$$\lambda_{n-k}(F(A_m)) \geq u'_m C_{L_m}' F(A_m) C_{L_m} u_m / \|C_{L_m} u\|^2 = u'_m F(A_m) u_m \geq \lambda_1(F(A_m)),$$

and that, using $\lambda_{n-k}(F(A_{m'})) - \lambda_1(F(A_{m'})) \rightarrow 0$ together with continuity of F at A , the upper and lower bounds in the previous display converge along m'' to $\lambda_1(F(A))$. It follows that $u'_* F(A) u_* = \lambda_1(F(A))$, and hence $u_* \in \text{Eig}(F(A), \lambda_1(F(A))) = \text{Eig}(A, \lambda_1(A))$, the equality following from equation (B.12). But from equation (B.11), we conclude that

$$\lambda_{\min}(M) < c = u'_m A_m u_m = u'_* A u_* = \lambda_1(A).$$

To arrive at a contradiction, it suffices to show that $\lambda_{\min}(M) = \lambda_{\min}(A)$. But (similar as argued above in the proof of (B.13)), this follows from equation (B.10), showing that $v'_m A_m v_m \rightarrow v'_* A v_* = \lambda_{\min}(M)$ along m'' , together with equation (B.15). \square

Proof of Proposition 3.6. We start with (1.): Let $\alpha \in (0, 1)$. Let X_m be a sequence of $n \times k$ -dimensional orthonormal matrices converging to some $Z \in \mathbb{R}^{n \times k}$ orthonormal, such that $e \notin \text{span}(X_m)$ holds, for every $m \in \mathbb{N}$, such that

$$T_{B(X_m), C_{X_m}}(e) - \lambda_1(B(X_m)) = e' C_{X_m}' B(X_m) C_{X_m} e / \|C_{X_m} e\|^2 - \lambda_1(B(X_m)) \rightarrow 0, \tag{B.16}$$

and such that $\liminf_{m \rightarrow \infty} \lambda_{n-k}(B(X_m)) - \lambda_1(B(X_m)) = \delta > 0$, where δ is a real number. Such a sequence exists as a consequence of Part 2 of Lemma B.2 (applied with $d = k$ and $v = e$). Without loss of generality, passing to a subsequence if necessary, we assume that $\lambda_{n-k}(B(X_m)) - \lambda_1(B(X_m)) > 0$ holds, for every $m \in \mathbb{N}$. Denote by κ_m the critical value $\kappa(\alpha)$ corresponding to $\Phi_{B(X_m), C_{X_m}, \kappa(\alpha)}$, cf. Lemma 2.3, and recall from that lemma that $\lambda_1(B(X_m)) < \kappa_m < \lambda_{n-k}(B(X_m))$ then holds as $\alpha \in (0, 1)$. Passing to a subsequence if necessary, we can assume that C_{X_m} converges to D_Z , say, an $(n - k) \times n$ matrix, the rows of which form an orthonormal basis of $\text{span}(Z)^\perp$. Recall the continuity property of F and that $B(X_m) = F(C_{X_m} M C_{X_m}')$. It follows that $B(X_m)$, $\lambda_1(B(X_m))$, and $\lambda_{n-k}(B(X_m))$ converge to $H := F(D_Z M D_Z')$, $b := \lambda_1(H)$, and $c := \lambda_{n-k}(H)$, respectively, with $c - b \geq \delta > 0$. Passing to another subsequence, if necessary, we can additionally achieve that $\kappa_m \rightarrow \kappa^*$, say. Obviously, $b \leq \kappa^* \leq c$ holds. We now argue that $b < \kappa^*$ must hold: By the definition of κ_m ,

$$\begin{aligned} \alpha &= P_{0,1,0}^{X_m} \left(\Phi_{B(X_m), C_{X_m}, \kappa(\alpha)} \right) = P_{0,1,0}^{X_m} \left(\Phi_{B(X_m), C_{X_m}, \kappa_m} \right) \\ &= P_{0,1,0}^{X_m} \left(\{y \in \mathbb{R}^n : T_{B(X_m)}(y) > \kappa_m\} \right). \end{aligned}$$

Denoting by G_m the cdf of the image measure $P_{0,1,0}^{X_m} \circ T_{B(X_m), C_{X_m}}$, this implies $1 - \alpha = G_m(\kappa_m)$. From Lemma B.4 of Preinerstorfer and Pötscher (2017), we obtain that the support of $P_{0,1,0}^{X_m} \circ T_{B(X_m), C_{X_m}}$ coincides with $[\lambda_1(B(X_m)), \lambda_{n-k}(B(X_m))]$, that G_m is a continuous function, and that G_m is strictly increasing on $[\lambda_1(B(X_m)), \lambda_{n-k}(B(X_m))]$. Hence, from $1 - \alpha \in (0, 1)$, it follows that $G_m^{-1}(1 - \alpha) = \kappa_m$, where G_m^{-1} denotes the quantile function corresponding to G_m . It is easy to see that G_m converges in distribution to the cdf G , say, of $P_{0,1,0}^Z \circ T_{H, D_Z}$, where the function $T_{H, D_Z} : \mathbb{R}^n \rightarrow \mathbb{R}$ is defined as

$$T_{H, D_Z}(y) = \begin{cases} y' D_Z' H D_Z y / \|D_Z y\|^2 & \text{if } y \notin \text{span}(Z), \\ \lambda_1(H) & \text{else.} \end{cases}$$

Again, Lemma B.4 of Preinerstorfer and Pötscher (2017) (with “ $B = H$ and $C_X = D_Z$ ”) shows that the support of G is $[b, c]$, that G is continuous (recall that $c - b \geq \delta > 0$), and that G is strictly increasing on $[b, c]$. This implies that the quantile function G^{-1} corresponding to G is continuous on $(0, 1)$, and that $G^{-1}(1 - \alpha) > b$. Using that $G_m \rightarrow G$ (in distribution), we conclude that the quantiles $\kappa_m = G_m^{-1}(1 - \alpha) \rightarrow G^{-1}(1 - \alpha) = \kappa^* > b$. Using equation (B.16), we can now conclude that there exists an $m_* \in \mathbb{N}$ such that $X_{m_*} =: X_*$ is of full column rank, such that $e \notin \text{span}(X_*)$, such that $\lambda_1(B(X_*)) < \lambda_{n-k}(B(X_*))$, and such that $T_{B(X_*), C_{X_*}}(e) < \kappa_{m_*}$ (with κ_{m_*} the critical value $\kappa(\alpha)$ corresponding to $\Phi_{B(X_*), C_{X_*}, \kappa(\alpha)}$ and $\alpha \in (0, 1)$). Theorem 3.2 establishes $X_* \in \mathcal{X}(\alpha; B)$.

We now prove (2.): Recall that X_* has full column rank and $e \notin \text{span}(X_*)$. We conclude that both statements (i) X is of full column rank and (ii) $e \notin \text{span}(X)$ hold for every X in an open set \mathcal{N} , say, containing X_* . We now claim that

$$T_{B(X), C_X}(e) < \bar{\kappa}(X); \text{ with } \bar{\kappa}(X) \in (\lambda_1(B(X)), \lambda_{n-k}(B(X))) \text{ s.t.}$$

$$P_{0,1,0}(\Phi_{B(X), C_X, \bar{\kappa}(X)}) = \alpha,$$

holds for every X in an open set $\mathcal{O} \subseteq \mathcal{N}$ containing X_* (that X_* satisfies the display was just established in the proof of (1.)). Arguing as in the proof of Part (1.), this claim and Theorem 3.2 (together with Lemma 2.3) would imply $\mathcal{O} \subseteq \mathcal{X}(\alpha; B)$, and we were done. To prove the claim, it suffices to verify that $T_{B(X), C_X}(e)$ and $\bar{\kappa}(X)$ as in the previous display are

(well-defined) continuous functions of X on a neighborhood of X_* . First, in order to ensure via Lemma 2.3 that a $\bar{\kappa}(X)$ as in the previous display uniquely exists on a neighborhood of X_* , we show that $\lambda_1(B(X)) < \lambda_{n-k}(B(X))$ holds on an open subset of \mathcal{N} containing X_* . Recalling that $\lambda_1(B(X_*)) < \lambda_{n-k}(B(X_*))$, and noting that the map $y \mapsto C_{X_*}y$ is a surjection of $\mathbb{R}^n \setminus \text{span}(X_*)$ to $\mathbb{R}^{n-k} \setminus \{0\}$, we conclude that there exist two vectors y_1 and y_2 in $\mathbb{R}^n \setminus \text{span}(X_*)$, and such that

$$\lambda_1(B(X_*)) = T_{B(X_*), C_{X_*}}(y_1) < T_{B(X_*), C_{X_*}}(y_2) = \lambda_{n-k}(B(X_*))$$

holds. From the additional continuity property in (2.), it follows that $y_1, y_2 \notin \text{span}(X)$ and $T_{B(X), C_X}(y_1) < T_{B(X), C_X}(y_2)$ hold on an open set $\mathcal{O}_1 \ni X_*$, say, such that $\mathcal{O}_1 \subseteq \mathcal{N}$, from which it follows that, for every $X \in \mathcal{O}_1$, we have $\lambda_1(B(X)) < \lambda_{n-k}(B(X))$. From $\mathcal{O}_1 \subseteq \mathcal{N}$, we conclude from Lemma 2.3 that a $\bar{\kappa}(X)$ satisfying the property to the right in penultimate display uniquely exists, for every $X \in \mathcal{O}_1$. Since $X \mapsto T_{B(X), C_X}(e)$ is continuous on $\mathcal{O}_1 \subseteq \mathcal{N}$ by assumption, it remains to verify that $X \mapsto \bar{\kappa}(X)$ is continuous on \mathcal{O}_1 . Lemma B.4 of Preinerstorfer and Pötscher (2017) and the definition of $\bar{\kappa}(X)$ show that, for $X \in \mathcal{O}_1$, we have $\bar{\kappa}(X) = F_X^{-1}(1 - \alpha)$, where F_X denotes the cdf of the image measure $P_{0,1,0} \circ T_{B(X), C_X}$. It is easy to see (using the additional continuity condition in (2.)) that the map $X \mapsto F_X$ is continuous on \mathcal{O}_1 (equipping the codomain with the topology of weak convergence). Furthermore, for every $X \in \mathcal{O}_1$, it holds (via Lemma B.4 in Preinerstorfer and Pötscher (2017)) that $P_{0,1,0} \circ T_{B(X), C_X}$ has support $[\lambda_1(B(X)), \lambda_{n-k}(B(X))]$ (which is nondegenerate), that the cdf F_X is continuous, and strictly increasing on $[\lambda_1(B(X)), \lambda_{n-k}(B(X))]$. Hence, for every $X \in \mathcal{O}_1$, the quantile function F_X^{-1} is continuous at $1 - \alpha \in (0, 1)$. Continuity of $X \mapsto \bar{\kappa}(X) = F_X^{-1}(1 - \alpha)$ on \mathcal{O}_1 follows. \square

Proof for the claim made in Remark 3.7. We verify that, for $B(X) = -(C_X \Sigma(\bar{\rho}) C'_X)^{-1}$, $\bar{\rho} \in (0, a)$, and every $z \in \mathbb{R}^n$, the function $X \mapsto T_{B(X), C_X}(z)$ is continuous at every $X \in \mathbb{R}^{n \times k}$ of full column rank such that $z \notin \text{span}(X)$. Fix $z \in \mathbb{R}^n$. Let X be of full column rank such that $z \notin \text{span}(X)$, and let X_m be a sequence converging to X . Eventually, X_m is of full column rank and satisfies $z \notin \text{span}(X_m)$, hence we may assume that this is the case for the whole sequence. We need to show that as $m \rightarrow \infty$ we have $T_{B(X_m), C_{X_m}}(z) \rightarrow T_{B(X), C_X}(z)$, or equivalently that

$$\frac{z' C'_{X_m} (C_{X_m} \Sigma(\bar{\rho}) C'_{X_m})^{-1} C_{X_m} z}{z' \Pi_{\text{span}(X_m)^\perp} z} \rightarrow \frac{z' C'_X (C_X \Sigma(\bar{\rho}) C'_X)^{-1} C_X z}{z' \Pi_{\text{span}(X)^\perp} z}.$$

Since X is of full column rank, $z' \Pi_{\text{span}(X_m)^\perp} z \rightarrow z' \Pi_{\text{span}(X)^\perp} z \neq 0$ obviously holds. For the numerators, let m' be an arbitrary subsequence of m , and choose m'' a subsequence of m' such that along m'' the sequence C_{X_m} converges to D , say. Note that D is necessarily orthonormal and $\text{span}(D) = \text{span}(X)^\perp$. Hence, along m'' , noting that $\Sigma(\bar{\rho})$ is positive definite by assumption, we have

$$z' C'_{X_m} (C_{X_m} \Sigma(\bar{\rho}) C'_{X_m})^{-1} C_{X_m} z \rightarrow z' D' (D \Sigma(\bar{\rho}) D')^{-1} D z.$$

Since $D = UC_X$ holds for an $(n - k) \times (n - k)$ orthonormal matrix U , say, it follows that

$$\begin{aligned} z' D' (D \Sigma(\bar{\rho}) D')^{-1} D z &= z' C'_X U' (UC_X \Sigma(\bar{\rho}) C'_X U')^{-1} UC_X z \\ &= z' C'_X (C_X \Sigma(\bar{\rho}) C'_X)^{-1} C_X z. \end{aligned}$$

Since the subsequence m' was arbitrary, we are done. \square

C. Proofs for Results in Section 4

Proof of Theorem 4.3. Denote by $\bar{P}_{(\beta,\gamma),\sigma,\rho}$ the distribution induced by (1), but where X is replaced by $\bar{X} = (X, e)$ (a matrix with column rank $k + 1 < n$), and where γ is the regression coefficient corresponding to e . Note also that, for every $\beta \in \mathbb{R}^k$, every $\sigma \in (0, \infty)$, and every $\rho \in [0, a)$, the measure $\bar{P}_{(\beta,0),\sigma,\rho}$ coincides with $P_{\beta,\sigma,\rho}$. An application of Corollary 2.22 in Preinerstorfer and Pötscher (2017) (recall that $\bar{\kappa}(\alpha) \in (\lambda_1(\bar{B}) < \lambda_{n-k-1}(\bar{B}))$ from the discussion preceding equation (10), and acting as if \bar{X} was the underlying design matrix), one then immediately obtains that, for every $\beta \in \mathbb{R}^k$, every $\sigma \in (0, \infty)$, and every $\gamma \in \mathbb{R}$, it holds that

$$0 < \lim_{\rho \rightarrow a} \bar{P}_{(\beta,\gamma),\sigma,\rho}(\bar{\Phi}_{\bar{B},\bar{\kappa}(\alpha)}) = \Pr(\bar{T}_{\bar{B}}(\Lambda \mathbf{G}) > \bar{\kappa}(\alpha)) < 1.$$

Setting $\gamma = 0$ then delivers the claim. □

Proof of Proposition 4.5. We proceed in three steps: The first step reduces the problem to an equivalent but simpler one by an invariance argument. The second step establishes a continuity property, and the last step establishes the equivalent formulation of the problem derived in the first step.

(1) By a simple G_X -invariance argument (recall A.1 and that $T_{C_X e e' C'_X}$ is G_X -invariant), it suffices to verify that, for every $\alpha \in (0, 1)$ and every $\varepsilon \in (0, \alpha)$, there exists a $c(\alpha, \varepsilon) \in (0, \kappa(\varepsilon))$ such that

$$\mathbb{E}_{0,1,0} \left[\min \left(\varphi_{\alpha-\varepsilon} + \mathbf{1}_{\Phi_{C_X e e' C'_X, c(\alpha,\varepsilon)}}, 1 \right) \right] = \alpha, \tag{C.1}$$

and such that, for every $c' \in (0, c(\alpha, \varepsilon))$, it holds that the supremum in (C.1) is greater than α .

(2) We claim that the nonincreasing function $g : \mathbb{R} \rightarrow \mathbb{R}$ defined via

$$c \mapsto \mathbb{E}_{0,1,0} \left[\min \left(\varphi_{\alpha-\varepsilon} + \mathbf{1}_{\Phi_{C_X e e' C'_X, c}}, 1 \right) \right]$$

is continuous. To verify this claim, let $c \in \mathbb{R}$, and let $c_m \rightarrow c$ be a real sequence. By the Dominated Convergence Theorem, to show that $g(c_m) \rightarrow g(c)$ holds, it is enough to verify

$$\lim_{m \rightarrow \infty} \left[\min \left(\varphi_{\alpha-\varepsilon}(y) + \mathbf{1}_{\Phi_{C_X e e' C'_X, c_m}}(y), 1 \right) \right] = \left[\min \left(\varphi_{\alpha-\varepsilon}(y) + \mathbf{1}_{\Phi_{C_X e e' C'_X, c}}(y), 1 \right) \right],$$

for $P_{0,1,0}$ -almost every $y \in \mathbb{R}^n$. To this end, it suffices to verify that

$$\lim_{m \rightarrow \infty} \mathbf{1}_{\Phi_{C_X e e' C'_X, c_m}}(y) = \mathbf{1}_{\Phi_{C_X e e' C'_X, c}}(y)$$

holds, for $P_{0,1,0}$ -almost every $y \in \mathbb{R}^n$. The statement in the previous display holds for every y such that $T_{C_X e e' C'_X}(y) \neq c$. The claim now follows from $P_{0,1,0}(\{y \in \mathbb{R}^n : T_{C_X e e' C'_X}(y) = c\}) = 0$, which can be obtained from Part 1 of Lemma B.4 in Preinerstorfer and Pötscher (2017) upon noting that $\lambda_1(C_X e e' C'_X) = 0$ (recall that $k < n - 1$) and that $0 < \|C_X e\|^2 = \lambda_{n-k}(C_X e e' C'_X)$ (the inequality following from $e \notin \text{span}(X)$).

(3) We now verify the statement around equation (C.1). Note that $\alpha - \varepsilon \leq g \leq 1$ (using A.1 and A.2 for the lower bound). Observe that $g(0) = 1$ follows from $1 \geq g(0) \geq P_{0,1,0}(\Phi_{C_X e e' C'_X, 0}) = 1$, the last equality following from Part 1 of Lemma B.4 in

Preinerstorfer and Pötscher (2017). Observe also that $g(\|C_X e\|^2) = \alpha - \varepsilon$ follows from $\alpha - \varepsilon \leq g(\|C_X e\|^2) \leq \alpha - \varepsilon + P_{0,1,0}(\Phi_{C_X e e' C'_X, \|C_X e\|^2}) = \alpha - \varepsilon$,

the last equality following again from Part 1 of Lemma B.4 in Preinerstorfer and Pötscher (2017). From these two observations, monotonicity of g and the continuity of g , it follows that $\{c \in \mathbb{R} : g(c) = \alpha\}$ is a closed interval contained in $(0, \|C_X e\|^2)$. Define $c(\alpha, \varepsilon)$ as the lower endpoint of this closed interval. Equation (C.1) and thus equation (14) follow. Furthermore, since $c(\alpha, \varepsilon)$ was defined as the lower endpoint, monotonicity of g implies that every $c' \in (0, c(\alpha, \varepsilon))$ must satisfy $g(c') > g(c(\alpha, \varepsilon)) = \alpha$. It remains to show that $c(\alpha, \varepsilon) \leq \kappa(\varepsilon)$. To this end, from what has already been established, it suffices to verify $g(\kappa(\varepsilon)) \leq \alpha$, which follows from

$$g(\kappa(\varepsilon)) = \mathbb{E}_{0,1,0} \left[\min \left(\varphi_{\alpha - \varepsilon} + \mathbf{1}_{\Phi_{C_X e e' C'_X, \kappa(\varepsilon)}}, 1 \right) \right] \\ \leq \mathbb{E}_{0,1,0}(\varphi_{\alpha - \varepsilon}) + P_{0,1,0}(\Phi_{C_X e e' C'_X, \kappa(\varepsilon)}) = \alpha.$$

□

Proof of Theorem 4.6. We start with the first claim in the theorem.

(1.) Let $\varepsilon \in (0, \alpha)$. Obviously

$$\varphi_{\alpha, \varepsilon}^* \geq \mathbf{1}_{\Phi_{C_X e e' C'_X, c(\alpha, \varepsilon)}}$$

which shows that, for every $\beta \in \mathbb{R}^k$, every $\sigma \in (0, \infty)$, and every $\rho \in [0, a)$, we have

$$\mathbb{E}_{\beta, \sigma, \rho}(\varphi_{\alpha, \varepsilon}^*) \geq P_{\beta, \sigma, \rho}(\Phi_{C_X e e' C'_X, c(\alpha, \varepsilon)}).$$

From the proof of Proposition 4.5, we know that

$$0 = \lambda_1(C_X e e' C'_X) < c(\alpha, \varepsilon) < \lambda_{n-k}(C_X e e' C'_X) = \|C_X e\|^2.$$

We can therefore use Lemma 2.3 (with $B = C_X e e' C'_X$) to conclude that $c(\alpha, \varepsilon) = \kappa(\alpha^*)$, for some $\alpha^* \in (0, 1)$, and apply Theorem 4.1 to conclude that, for every $\beta \in \mathbb{R}^k$ and every $\sigma \in (0, \infty)$, we have $\lim_{\rho \rightarrow a} P_{\beta, \sigma, \rho}(\Phi_{C_X e e' C'_X, c(\alpha, \varepsilon)}) = 1$, which together with the lower bound in the previous display proves the first claim. We move on to the second claim in the theorem.

(2.) Using G_X -invariance of $\varphi_{\alpha, \varepsilon}^*$ (for every $\varepsilon \in (0, \alpha)$) and of φ_α , together with $\|\Sigma(\rho)\| > 0$, for every $\rho \in [0, a)$, it suffices to verify that

$$\lim_{\varepsilon \rightarrow 0^+} \sup_{\rho \in A} |\mathbb{E}_{0, \|\Sigma(\rho)\|^{-1/2, \rho}}(\varphi_{\alpha, \varepsilon}^*) - \mathbb{E}_{0, \|\Sigma(\rho)\|^{-1/2, \rho}}(\varphi_\alpha)| = 0.$$

Let $\varepsilon_m \rightarrow 0$ be a sequence in $(0, \alpha)$ and let ρ_m be a sequence in A . For convenience, set $\sigma_m := \|\Sigma(\rho_m)\|^{-1/2}$. We verify that

$$|\mathbb{E}_{0, \sigma_m, \rho_m}(\varphi_{\alpha, \varepsilon_m}^*) - \mathbb{E}_{0, \sigma_m, \rho_m}(\varphi_\alpha)| = |\mathbb{E}_{0, \sigma_m, \rho_m}(\varphi_{\alpha, \varepsilon_m}^* - \varphi_\alpha)| \rightarrow 0. \tag{C.2}$$

Let m' be an arbitrary subsequence of m . By compactness of the unit sphere in $\mathbb{R}^{n \times n}$, we can choose a subsequence m'' of m' along which $\|\Sigma(\rho_m)\|^{-1} \Sigma(\rho_m)$ converges to a symmetric matrix Γ , say, which due to the additional assumption on the set A is positive definite. It follows from Scheffé’s lemma that along m'' the sequence P_{0, σ_m, ρ_m} (i.e., the Gaussian probability measure with mean 0 and covariance matrix $\|\Sigma(\rho_m)\|^{-1} \Sigma(\rho_m)$) converges in

total variation distance to Q , a Gaussian probability measure with mean 0 and covariance matrix Γ . Clearly

$$|\mathbb{E}_{0, \sigma_m, \rho_m}(\varphi_{\alpha, \varepsilon_m}^* - \varphi_\alpha)| \leq \mathbb{E}_{0, \sigma_m, \rho_m}(|\varphi_{\alpha, \varepsilon_m}^* - \varphi_\alpha|).$$

By, e.g., Lemma 2.3 in Strasser (1985) and since $|\varphi_{\alpha, \varepsilon_m}^* - \varphi_\alpha|$ is a sequence of tests, it follows from the total variation convergence of the sequence P_{0, σ_m, ρ_m} to Q along m'' , that along m'' we have

$$|\mathbb{E}_{0, \sigma_m, \rho_m}(|\varphi_{\alpha, \varepsilon_m}^* - \varphi_\alpha|) - \mathbb{E}_Q(|\varphi_{\alpha, \varepsilon_m}^* - \varphi_\alpha|)| \rightarrow 0,$$

where \mathbb{E}_Q denotes expectation w.r.t. Q . We now claim that

$$\mathbb{E}_Q(|\varphi_{\alpha, \varepsilon_m}^* - \varphi_\alpha|) \rightarrow 0. \tag{C.3}$$

This claim, if true, then implies equation (C.2) as the subsequence m' we started with was arbitrary.

Instead of (C.3), we first show that $\mathbb{E}_{0, 1, 0}(|\varphi_{\alpha, \varepsilon_m}^* - \varphi_\alpha|) \rightarrow 0$. To this end, write

$$\varphi_{\alpha, \varepsilon_m}^* - \varphi_\alpha = [\varphi_{\alpha - \varepsilon_m} - \varphi_\alpha] + (1 - \varphi_{\alpha - \varepsilon_m}(y)) \mathbf{1}_{\Phi_{C_X e' C_X' c(\alpha, \varepsilon_m)}}. \tag{C.4}$$

From A.3 and the Dominated Convergence Theorem, we obtain $\mathbb{E}_{0, 1, 0}[|\varphi_{\alpha - \varepsilon_m} - \varphi_\alpha|] \rightarrow 0$. It remains to show that $\mathbb{E}_{0, 1, 0}(\psi_m) \rightarrow 0$, for

$$\psi_m := (1 - \varphi_{\alpha - \varepsilon_m}(y)) \mathbf{1}_{\Phi_{C_X e' C_X' c(\alpha, \varepsilon_m)}} \geq 0.$$

By construction and A.2, we have $\mathbb{E}_{0, 1, 0}(\varphi_{\alpha, \varepsilon_m}^*) = \alpha = \mathbb{E}_{0, 1, 0}(\varphi_\alpha)$. Therefore, (C.4) shows that $-\mathbb{E}_{0, 1, 0}[\varphi_{\alpha - \varepsilon_m} - \varphi_\alpha] = \mathbb{E}_{0, 1, 0}(\psi_m)$. Hence, $\mathbb{E}_{0, 1, 0}(\psi_m) \rightarrow 0$ follows from $\mathbb{E}_{0, 1, 0}[\varphi_{\alpha - \varepsilon_m} - \varphi_\alpha] \rightarrow 0$. This shows that $\mathbb{E}_{0, 1, 0}(|\varphi_{\alpha, \varepsilon_m}^* - \varphi_\alpha|) \rightarrow 0$.

To show (C.3), let m^* be a subsequence of m . Since $\mathbb{E}_{0, 1, 0}(|\varphi_{\alpha, \varepsilon_{m^*}}^* - \varphi_\alpha|) \rightarrow 0$, there exists a subsequence m^{**} of m^* , and a set N such that $P_{0, 1, 0}(N) = 0$, and such that, for every $y \in \mathbb{R}^n \setminus N$, it holds that $|\varphi_{\alpha, \varepsilon_{m^{**}}}^*(y) - \varphi_\alpha(y)| \rightarrow 0$ (cf., e.g., Theorem 3.12 in Rudin (1987)). From positive definiteness of Γ , it follows that $Q(N) = 0$, and (by the Dominated Convergence Theorem) that $\lim_{m^{**} \rightarrow \infty} \mathbb{E}_Q(|\varphi_{\alpha, \varepsilon_{m^{**}}}^* - \varphi_\alpha|) = 0$. \square

REFERENCES

Davies, R.B. (1980) Algorithm AS 155: The distribution of a linear combination of χ^2 random variables. *Journal of the Royal Statistical Society. Series C (Applied Statistics)* 29(3), 323–333.

Fan, J., Y. Liao, & J. Yao (2015) Power enhancement in high-dimensional cross-sectional tests. *Econometrica* 83(4), 1497–1541.

Horn, R.A. & C.R. Johnson (1985) *Matrix Analysis*. Cambridge: Cambridge University Press.

Kelly, M. (2019) The standard errors of persistence. *CEPR Discussion paper no. DP13783*.

King, M.L. & G.H. Hillier (1985) Locally best invariant tests of the error covariance matrix of the linear regression model. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* 47, 98–102.

Kleiber, C. & W. Krämer (2005) Finite-sample power of the Durbin–Watson test against fractionally integrated disturbances. *Econometrics Journal* 8(3), 406–417.

Kock, A.B. & D. Preinerstorfer (2019) Power in high-dimensional testing problems. *Econometrica* 87(3), 1055–1069.

Krämer, W. (1985) The power of the Durbin–Watson test for regressions without an intercept. *Journal of Econometrics* 28(3), 363–370.

- Krämer, W. (2005) Finite sample power of Cliff–Ord-type tests for spatial disturbance correlation in linear regression. *Journal of Statistical Planning and Inference* 128(2), 489–496.
- Krämer, W. & H. Zeisel (1990) Finite sample power of linear regression autocorrelation tests. *Journal of Econometrics* 43(3), 363–372.
- Leeb, H. & B.M. Pötscher (2017) Testing in the presence of nuisance parameters: Some comments on tests post-model-selection and random critical values. In S.E. Ahmed (ed.), *Big and Complex Data Analysis: Methodologies and Applications*, pp. 69–82. Springer.
- Lehmann, E.L. & J.P. Romano (2005) *Testing Statistical Hypotheses, 3rd Edition*. Springer.
- Löbus, J.-U. & L. Ritter (2000) The limiting power of the Durbin–Watson test. *Communications in Statistics: Theory and Methods* 29(12), 2665–2676.
- Martellosio, F. (2010) Power properties of invariant tests for spatial autocorrelation in linear regression. *Econometric Theory* 26(1), 152–186.
- Martellosio, F. (2012) Testing for spatial autocorrelation: The regressors that make the power disappear. *Econometric Reviews* 31(2), 215–240.
- Preinerstorfer, D. (2017) Finite sample properties of tests based on prewhitened nonparametric covariance estimators. *Electronic Journal of Statistics* 11(1), 2097–2167.
- Preinerstorfer, D. & B.M. Pötscher (2016) On size and power of heteroskedasticity and autocorrelation robust tests. *Econometric Theory* 32(2), 261–358.
- Preinerstorfer, D. & B.M. Pötscher (2017) On the power of invariant tests for hypotheses on a covariance matrix. *Econometric Theory* 33(1), 1–68.
- Rudin, W. (1987) *Real and Complex Analysis*. McGraw-Hill Education.
- Strasser, H. (1985) *Mathematical Theory of Statistics: Statistical Experiments and Asymptotic Decision Theory*. Walter de Gruyter.
- Zeisel, H. (1989) On the power of the Durbin–Watson test under high autocorrelation. *Communications in Statistics: Theory and Methods* 18(10), 3907–3916.