

A Service of



Leibniz-Informationszentrum Wirtschaft Leibniz Information Centre for Economics

Muris, Chris; Wacker, Konstantin M.

Working Paper Estimating interaction effects with panel data

GLO Discussion Paper, No. 1583

Provided in Cooperation with: Global Labor Organization (GLO)

Suggested Citation: Muris, Chris; Wacker, Konstantin M. (2025) : Estimating interaction effects with panel data, GLO Discussion Paper, No. 1583, Global Labor Organization (GLO), Essen

This Version is available at: https://hdl.handle.net/10419/313870

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



WWW.ECONSTOR.EU

Estimating interaction effects with panel data^{*}

Chris Muris[†] Konstantin M. Wacker^{\ddagger §}

March 24, 2025

Abstract

This paper analyzes how interaction effects can be consistently estimated under economically plausible assumptions in linear panel models with a fixed *T*dimension. We advocate for a *correlated interaction term estimator* (CITE) and show that it is consistent under conditions that are not sufficient for consistency of the interaction term estimator that is most common in applied econometric work. Our paper discusses the empirical content of these conditions, shows that standard inference procedures can be applied to CITE, and analyzes consistency, relative efficiency, inference, and their finite sample properties in a simulation study. In an empirical application, we test whether labor displacement effects of robots are stronger in countries at higher income levels. The results are in line with our theoretical and simulation results and indicate that standard interaction term estimation underestimates the importance of a country's income level in the relationship between robots and employment and may prematurely reject a null hypothesis about interaction effects in the presence of misspecification.

Keywords: panel data, interaction effects, correlated random coefficients, robots.

JEL codes: C13, C23, C33, J23, O30.

^{*}This paper benefited from discussions at presentations at the Universities of Bristol, Göttingen, Groningen, Mainz, WU Vienna, the World Bank, the European Commission's Joint Research Centre, and the Southern Economic Association meetings. We would like to thank the respective participants for their feedback and inputs; comments from Kirill Borusyak, Irene Botosaru, Peter Egger, Michaela Kesina, Adam Lavecchia, Krishna Pendakur, James Powell, Nathanael Vellekoop, Frank Windmeijer, and Jonathan Zhang further helped to significantly improve our paper. We are additionally grateful to Mathieu Couttenier, Mathias Thoenig, and to the Swiss BfS for providing data we used for an empirical application that entered an earlier working paper version of our paper (Muris & Wacker, 2022), and to Eline Koopman for excellent research assistance.

[†]Department of Economics, McMaster University, muerisc@mcmaster.ca

[‡]Department of Global Economics & Management, University of Groningen, k.m.wacker@rug.nl [§]GLO Fellow (Global Labor Organization)

1 Introduction

In many econometric applications, the relationship between an outcome variable Y and a regressor X depends on a so-called *interaction variable* H. For example, an earlier macroeconomic literature postulated that the effect of foreign capital inflows X on economic growth Y depends on some host-country fundamentals H, such as institutional quality, financial development, or a sufficiently educated labor force.¹ In micro data, the effect of new technologies on firms' productivity may depend on firm size, reflecting scale effects. And how a worker responds to a labor market reform may depend on age, reflecting that older individuals value leisure higher and income less in their labor supply choice than younger workers.

Panel data have at least two dimensions, i and t, which provides us with opportunities to estimate interaction effects that uni-dimensional data do not have. In particular, panel data give us more flexibility to model heterogeneities in the relationship between Y and X, on top of variation in H. Yet, those opportunities are rarely exploited in applied econometric analysis, which possibly reflects that theoretical work to date has largely ignored the benefits of modeling heterogeneities across panel units for the estimation of interaction effects in panel data.

Consider the relationship between employment Y and robotization X. We may hypothesize that this relationship is more negative in countries i with a higher wage level H_i because firms experience higher pressure to save labor costs. The standard approach to estimating such an interaction effect in a linear model with panel data is to regress employment Y_{it} on robotization X_{it} , and its interaction with countries' wage levels H_i , together unit-specific fixed effects α_i and possibly additional variables in Z_{it} , where t may index time, industries, firms etc.:

$$Y_{it} = \alpha_i + X_{it}\beta + (X_{it}H_i)\kappa + Z'_{it}\gamma + \nu_{it}.$$
(1)

The conventional approach to estimate interaction effects in a linear model with panel data seems to be least squares based on equation (1).² We refer to this approach as the

¹See, for example, Alfaro, Chanda, Kalemli-Ozcan, and Sayek (2004, 2010); Borensztein, Gregorio, and Lee (1998); Burnside and Dollar (2000); Durham (2004). Economically, this literature assumes that foreign capital X and interaction variables H complement each other in the production of output Y.

²See, for example, Alsan and Goldin (2019); Amiti and Konings (2007); Berman, Martin, and Mayer (2012); Bloom, Draca, and Van Reenen (2016); Bloom, Sadun, and Van Reenen (2012); Burnside and Dollar (2000); Duflo, Dupas, and Kremer (2011); Epifani and Gancia (2009); Herrera, Ordoñez, and

interaction term estimator (ITE). The partial effect of X (robots) on Y (employment) linearly depends on H (wage level) in this model, as is easily verified by taking the partial derivative $\partial Y/\partial X = \beta + H\kappa$.

Our paper focuses on the interaction effect κ and how it can be consistently estimated under economically plausible assumptions with panel data that have a short Tdimension. Standard linear regression assumptions in application to equation (1) require the error term ν to be independent of the regressors X, XH, Z. This assumption is easily violated in many interaction models: the economic reasons for effect heterogeneity in the relationship between X and Y are extensive and rarely independent of the regressors. Recall the example of robots and employment. In a technology-minded country, firms t may find various ways to substitute workers' tasks in the production process by robots. The partial relationship between robots X and employment Y will hence be more negative in those countries than in less tech-minded countries because a robot can replace more workers in the former. Since this aspect is unmodeled in equation (1), and tech-mindedness is not directly observable, this heterogeneity across countries will end up in the error term ν . And since tech-mindedness likely gives rise to higher robotization X in the first place, we observe a standard endogeneity bias: ν is correlated with X and XH in equation (1).

Panel data models for interaction effects can be deceptive because the inclusion of the additive fixed effects α_i may lull the researcher into a false sense of having controlled for unobserved heterogeneity. This is not the case because in the context of interaction effects we are not concerned about (additive) heterogeneity in levels of Y and X, which are indeed absorbed through α_i , but about *effect heterogeneity* in the relationship between X and Y. Our research shows that a common source of such unobserved effect heterogeneity can be easily addressed with panel data. Our exposition focuses on the simple case where the interaction variable H only varies in the i dimension since this is plausibly the key dimension giving rise to heterogeneity in the partial effect of X on Y. An extension to interaction variables that vary in the it dimension and an application to higher-dimensionality panel data is provided in an earlier working paper version of this paper (Muris & Wacker, 2022) and follows the same rationale: additive fixed effects do not control for important aspects of heterogeneity.

The key contribution of our paper is to advocate for a different approach to the estimation of interaction effects with panel data and to analyze the exogeneity conditions

Trebesch (2020); List and Sturm (2006); Manacorda and Tesei (2020); Shambaugh (2004); Spilimbergo (2009); Storeygard (2016).

that are sufficient for consistent estimation of κ . We therefore develop a framework of correlated unobserved effect heterogeneity that explicitly includes unit-specific effects β_i as parameters in the relationship between X and Y (see esp. Chamberlain, 1992). Building on this framework, we propose a *correlated interaction term estimator* (CITE) that obtains an estimate for the interaction term coefficient κ by first running the regression $Y_{it} = \alpha_i + X_{it}\beta_i + Z'_{it}\gamma + U_{it}$ and subsequently projecting the unit-specific effects $\hat{\beta}_i$ onto H_i in a second-step regression.³ Since this second-step regression of $\hat{\beta}_i$ onto H_i is a simple cross-section OLS regression, it is easily to implement and can be assessed with the extensive, well-known statistical toolkit that has been developed for OLS regressions.⁴

A key result of the rigorous comparison of ITE and CITE that we provide in this paper is that the exogeneity restrictions that guarantee the consistency of CITE for κ are not sufficient for the consistency of ITE. On top of conditions that apply to both estimators, CITE requires unobserved effect heterogeneity to be uncorrelated with the interaction variable H, while ITE requires this effect heterogeneity to be uncorrelated with all right-hand side regressors H, X, Z for consistent estimation of κ . This suggests that researchers need strong exogeneity arguments when estimating interaction terms with ITE in panel data.

One plausible concern about our newly proposed CITE is that its superior bias protection comes at the cost of efficiency. By treating the unit-specific effects β_i as parameters that need to be estimated, one may expect a higher variability of CITE for κ , compared to ITE. We present simulation results demonstrating that this is not necessarily the case when considering different Monte Carlo designs. There is hence no obvious penalty for using CITE in terms of efficiency, while its advantages in terms of consistency are clear.

Another advantage of CITE is that we can demonstrate that standard inference procedures, available in common software, can be applied under economically plausible

³Such an approach has occasionally been used in applied work. See particularly Couttenier, Petrencu, Rohner, and Thoenig (2019) and MaCurdy (1981). Giesselmann and Schmidt-Catran (2020) discuss CITE, but make a case for a double-demeaned estimator. Alternatively, running unit-by-unit regressions to recover common parameters is a strategy that has been used in different panel data settings, see e.g. Fernández-Val and Lee (2013). In any case, we are not aware of existing work that demonstrates that ITE and CITE are distinct and that analyzes their asymptotic properties.

⁴If one is interested in time-varying interaction variables, those can be considered part of Z and the second step can be omitted. Unobserved effect heterogeneity across panel units is directly absorbed in $\hat{\beta}_i$ in this case. See an earlier working paper version of our paper for a more explicit treatment of such time-varying interactions (Muris & Wacker, 2022).

assumptions about the structure of error terms. The finite sample performance of these inference procedures is also demonstrated in our simulation section. Finally, CITE does not require a large number of time periods. Our paper develops theoretical consistency results under fixed-T and our simulations suggest that relative efficiency is often similar to ITE around T = 4, depending on the simulation design. The fact that we can prove consistency of CITE for fixed-T also sets our paper apart from the literature on heterogeneity in large-T panels.⁵

Related literature

Our framework builds on the literature on correlated random coefficient models (see Arellano & Bonhomme, 2012; Chamberlain, 1992; Graham & Powell, 2012; Laage, 2024; Sasaki & Ura, 2021). Specifically, we build on Corollary 1 in Arellano and Bonhomme (2012) to study the estimation of interaction effects in panel models. We adapt and extend the theoretical results in this literature to advocate for CITE over ITE.

Models with correlated coefficient heterogeneity have previously been used to study the identification and estimation of average treatment effects in various settings, see for example Wooldridge (2005), Murtazashvili and Wooldridge (2008), de Chaisemartin and D'Haultfœuille (2022); Verdier (2020), Słoczyński (2022), de Chaisemartin and Lei (2023), Winkelmann (2024), de Chaisemartin and D'Haultfœuille (2024). This literature is mostly concerned with the identification and estimation of (convex combinations of) average coefficients, partial effects, and treatment effects, from cross-sectional and panel data. In contrast, we focus on the estimation of *interaction effects* from panel data in the presence of unobserved effect heterogeneity.

Organization and preview from an applied perspective

Section 2 introduces the unobservable effect heterogeneity model. Section 3 formally defines the two estimators. Section 4 presents our main results about the consistency of CITE. Section 5 discusses inference. Section 6 contains a Monte Carlo simulation study. Section 7 contains an empirical application where we investigate whether worker displacement effects of robots are stronger in countries at higher income levels. Section 8

⁵Balli and Sørensen (2013) observe that unobserved effect heterogeneity can lead to inconsistency in the ITE, and suggest that "if the time-series dimension of the data is large, one may directly allow for country-varying slopes". We show that such slopes should generally be preferred regardless of the time-series dimension.

concludes and discusses some avenues for future research.

In sections 2 - 5, which focus on theory, we have mostly opted for a notation of within-transformed stacked variables to facilitate the theoretical exposition. Applied readers who are at less comfort with this notation, or less interested in the details of our theoretical results, are suggested to focus on the following aspects: the part of section 2 prior to notational definitions to understand the unobservable effect heterogeneity model we have in mind and the first two paragraphs of section 3 for an informal definition of ITE and CITE.

The crucial exogeneity conditions required for consistent estimation of interaction terms in panel data can be found in Assumptions 6-8 in Section 4. It is particularly instructive to compare Assumption 7, which is sufficient for consistency of CITE, to the stronger Assumption 8, which is required for consistency of ITE. Remark 2 compares both assumptions with an example from our application.

The essence of section 5 for applied researchers is that conventional heteroskedasticityrobust standard errors in the second-stage projection of $\hat{\beta}_i$ onto H_i are sufficient for valid inference under CITE. Section 6 is accessible without delving into technical or notational details, but can also be skipped if one is already convinced about the advantages of CITE in terms of consistency, relative efficiency, inference, and their finite sample properties. The application section 7 will be most instructive for applied readers to understand the implementation of our approach in practice.

2 Model

We are interested in estimating *interaction effects* in static linear panel models with a short T dimension. In our framework, the effect of X on Y linearly depends on observables H and on additional *unobservable* sources of effect heterogeneity. We model this unobserved heterogeneity in the effect of X on Y as follows:

$$Y_{it} = \alpha_i + X_{it}\beta_i + Z'_{it}\gamma + U_{it}, \qquad i = 1, \cdots, n, \quad t = 1, \cdots, T, \quad (2)$$

$$\beta_i = \kappa_0 + \kappa_1 H_{1i} + \kappa_2 H_{2i} + \dots + \epsilon_i = H'_i \kappa + \epsilon_i, \qquad \qquad i = 1, \cdots, n, \qquad (3)$$

where i indexes cross-sectional units, and t indexes time periods (or other panel dimensions). The beginning of Section 3 clarifies how the standard panel interaction term model relates to this framework.

The *outcome equation* (2) of our framework describes how a dependent variable

 $Y_{it} \in \mathbb{R}$ responds to a change in a regressor $X_{it} \in \mathbb{R}$,⁶ allowing for additive fixed effects $\alpha_i \in \mathbb{R}$, control variables $Z_{it} \in \mathbb{R}^{K_z}$, and an error term $U_{it} \in \mathbb{R}$. It is important to note that the control variables Z_{it} may include interaction terms with *it* variation. As an earlier working paper version of our paper shows more thoroughly, CITE poses less restrictive exogeneity conditions for their consistency as well (Muris & Wacker, 2022).

The heterogeneity equation (3) of our framework relates effect heterogeneity β_i to observed, time-invariant interaction variables $H_i \in \mathbb{R}^{K_h}$. The associated interaction term coefficient $\kappa \in \mathbb{R}^{K_h}$ is the main object of interest in our paper.⁷ The effect of X on Y may additionally vary, even when holding H_i constant, because of unobserved effect heterogeneity $\epsilon_i \in \mathbb{R}$. The relationship between unobserved effect heterogeneity and (X, H) plays an important role in our analysis (in a similar way as the relationship between unobserved heterogeneity α_i and X is important in the conventional additive fixed effect model).

Remark 1. The index t need not refer to time, but can refer to students t for a given classroom i, counties t within a given state i, employees t within a given firm i, etc. The empirical application in Section 7 provides an example and section 6.2 in an earlier working paper version illustrates how our framework applies to higher-dimensional panel data (Muris & Wacker, 2022).

Some additional notation simplifies the exposition in the remainder of the paper. First, estimation will be based on the within-transformed outcome equation:

$$\widetilde{Y}_{it} = \widetilde{X}_{it}\beta_i + \widetilde{Z}'_{it}\gamma + \widetilde{U}_{it}, \qquad (4)$$

where

$$\widetilde{Y}_{it} = Y_{it} - \frac{1}{T} \sum_{s=1}^{T} Y_{is}, \quad \widetilde{X}_{it} = X_{it} - \frac{1}{T} \sum_{s=1}^{T} X_{is},$$

and \widetilde{Z}_{it} and \widetilde{U}_{it} are defined analogously.

⁶Our analysis is easily generalized to vector-valued X_{it} .

⁷Note that we obtain the constant κ_0 in the heterogeneity equation (3) by setting the first element of all H_i equal to 1 and that this constant can be interpreted as the average β_i across panel units, conditional on all other elements of H_i .

Second, the reduced form of the within-transformed equations (2)-(3) is

$$\widetilde{Y}_{it} = \widetilde{X}_{it}(H'_i\kappa + \epsilon_i) + \widetilde{Z}'_{it}\gamma + \widetilde{U}_{it}$$

$$= \underbrace{\widetilde{X}_{it}H'_i}_{\Psi'_{it}}\kappa + \widetilde{Z}'_{it}\gamma + \underbrace{\left(\widetilde{U}_{it} + \widetilde{X}_{it}\epsilon_i\right)}_{V_{it}},$$
(5)

with an error term V_{it} that depends on the unobserved effect heterogeneity ϵ_i .

Third, we collect the within-transformed outcome equation (4) across $t = 1, \dots, T$ for a given *i* to obtain

$$\widetilde{Y}_i = \widetilde{X}_i \beta_i + \widetilde{Z}_i \gamma + \widetilde{U}_i, \tag{6}$$

where

$$\widetilde{X}_i = (\widetilde{X}_{i1}, \cdots, \widetilde{X}_{iT}) \in \mathbb{R}^T,$$

and $\widetilde{Y}_i \in \mathbb{R}^T$, \widetilde{X}_i , $\widetilde{Z}_i \in \mathbb{R}^{T \times K_z}$, and $\widetilde{U}_i \in \mathbb{R}^T$. Accordingly, we obtain

$$\widetilde{Y}_i = \Psi_i \kappa + \widetilde{Z}_i \gamma + V_i, \tag{7}$$

where $\Psi_i = \widetilde{X}_i H'_i \in \mathbb{R}^{T \times K_h}$ and $V_i = \widetilde{U}_i + \widetilde{X}_i \epsilon_i \in \mathbb{R}^T$ are stacked versions of Ψ_{it} and V_{it} , respectively.

3 Estimators

We consider two estimators for the interaction term coefficient κ : the interaction term estimator (ITE) and the correlated interaction term estimator (CITE). The ITE is a fixed effects regression⁸ of Y on the interaction term $\Psi = XH$ and Z, as suggested by the reduced form (5). It is the default approach in applied research when the effect of X is expected to depend on H: "just add an interaction term". This standard ITE panel interaction model is usually written along the lines $Y_{it} = \alpha_i + X_{it}\beta + X_{it}H_{1i}\kappa_1 + Z'_{it}\gamma + \nu_{it}$ in many applied papers, which is a special case of our framework that one obtains by substituting equation (3) into equation (2), imposing $\epsilon_i = 0 \forall i$, and setting $\beta = \kappa_0$. A formal definition of ITE is provided in Section 3.2.

The CITE, defined formally in Section 3.1, is a two-step estimator. The first step is a fixed effects regression of Y on Z and interactions of X with dummy variables for each

 $^{^{8}\}mathrm{By}$ "fixed effects regression" we mean a regression that includes a dummy variable for each i, i.e. the within estimator.

i: $Y_{it} = \alpha_i + X_{it}\beta_i + Z'_{it}\gamma + U_{it}$. The coefficients on the dummy variable interactions, $\hat{\beta}_i$, are individual-specific effects of X on Y that capture unobserved effect heterogeneity. The second step is a regression of $\hat{\beta}_i$ on H. CITE hence directly follows the structure in equations (2)–(3): the first-step regression is based on the outcome equation (2), and the second-step regression is based on heterogeneity equation (3). Practical implementation of CITE in standard software is simple: it just requires interacting X with the unit-specific dummy variables, saving the estimated coefficients, and regressing them on H in a second step.

The remainder of this manuscript assumes a random sample is available.

Assumption 1 (Random sampling). For each $i = 1, \dots, n$, the observed data is

$$W_i = (Y_{i1}, \cdots, Y_{iT}, X_{i1}, \cdots, X_{iT}, Z_{i1}, \cdots, Z_{iT}, H_i),$$

generated by equations (2)–(3). The random sequence $\{W_i\}_{i=1}^n$ is independent and identically distributed.

3.1 Correlated interaction term estimator

Sufficient variation in \widetilde{X}_i for all *i* is required for the CITE to be well-defined.

Assumption 2. There exists an h > 0 such that

$$\inf_{i} \left(\sum_{t=1}^{T} \widetilde{X}_{it}^{2} \right) \ge h.$$

This guarantees that the least squares estimators for the (β_i) are well-defined, which avoids the identification issues addressed in Graham and Powell (2012) and Arellano and Bonhomme (2012). If Assumption 2 does not hold for all *i*, our analysis applies to the subpopulation for which it does hold.

To obtain the within-transformed data, we define the residual maker matrix

$$M_i = I_T - \widetilde{X}_i \left(\widetilde{X}'_i \widetilde{X}_i \right)^{-1} \widetilde{X}'_i,$$

where I_T is the $T \times T$ identity matrix. Premultiplication by M_i obtains

$$M_{i}\widetilde{Y}_{i} = M_{i}\widetilde{X}_{i}\beta_{i} + M_{i}\widetilde{Z}_{i}\gamma + M_{i}\widetilde{U}_{i}$$
$$= M_{i}\widetilde{Z}_{i}\gamma + M_{i}\widetilde{U}_{i}, \qquad (8)$$

and estimation of γ can be based on equation (8).

Assumption 3. The matrix $E\left(\widetilde{Z}'_i M_i \widetilde{Z}_i\right)$ is invertible.

Assumptions 1-3 guarantee that the following is well-defined for large n.

Definition 1. CITE for γ is given by

$$\widehat{\gamma}_n = \left(\sum_{i=1}^n \widetilde{Z}'_i M_i \widetilde{Z}_i\right)^{-1} \sum_{i=1}^n \widetilde{Z}'_i M_i \widetilde{Y}_i.$$

To define the CITE for κ , we also need sufficient variation in H_i .

Assumption 4. The matrix $E(H'_iH_i)$ is invertible.

For each i, the estimator for the individual-specific effect of X on Y is

$$\widehat{\beta}_i = \left(\widetilde{X}_i'\widetilde{X}_i\right)^{-1} X_i' \left(Y_i - Z_i\widehat{\gamma}_n\right).$$

The CITE for κ is obtained from a regression of $\hat{\beta}_i$ on H_i .

Definition 2. The CITE for κ is

$$\widehat{\kappa}_n = \left(\sum_{i=1}^n H'_i H_i\right)^{-1} \sum_{i=1}^n H'_i \widehat{\beta}_i.$$

3.2 Interaction term estimator

The ITE is based on equation (7). To formally define it, we project out the fixed effects α_i using the within transformation. For the ITE to be well-defined, we require the following 'no multicollinearity' condition:

Assumption 5. The matrix $E\left([\Psi_i \ \widetilde{Z}_i]'[\Psi_i \ \widetilde{Z}_i]\right)$ is positive definite.

Definition 3. Given Assumption 5, the ITE for $\theta = (\kappa, \gamma)$ is well-defined:

$$\check{\theta}_n = \left(\sum_{i=1}^n [\Psi_i \ \widetilde{Z}_i]' [\Psi_i \ \widetilde{Z}_i]\right)^{-1} \sum_{i=1}^n [\Psi_i \ \widetilde{Z}_i]' \widetilde{Y}_i$$

$$= (\check{\kappa}_n, \check{\gamma}_n).$$

4 Consistency

We now show that a set of exogeneity conditions that is sufficient for consistency of CITE is not sufficient for consistency of ITE. Throughout, we maintain the following strict exogeneity assumption.

Assumption 6 (Strict exogeneity). $E(U_i|X_i, Z_i, H_i) = 0.$

This assumption is similar to the strict exogeneity assumption that is standard in the literature on correlated random coefficient panel models,⁹ and in textbook treatments of fixed-*T* linear panel models with additive fixed effects. If one is only interested in γ , and not in κ , Assumption 6 can be relaxed for CITE (but not for ITE) to $E(U_i|X_i, Z_i) = 0$.

4.1 CITE

It is sufficient for the consistency of CITE that the unobserved effect heterogeneity ϵ_i is orthogonal to the time-invariant interaction variable H_i .

Assumption 7 (Exogeneity, ϵ). $E(\epsilon_i | H_i) = 0$.

This assumption is easily understood when recalling that the second step of CITE is a cross-sectional OLS regression: it is equivalent to the textbook OLS exogeneity condition that requires the regressor H_i to be orthogonal to the error term ϵ_i .

To understand the advantages of CITE in applied econometric work, it is crucial to emphasize that Assumption 7 does not involve the regressors (X, Z). We discuss this advantage in contrast with ITE, including an example, in Remark 2 at the end of Section 4.

Theorem 1. (i) If Assumptions 1–3 and 6 hold, and if $E ||Z_i||^2$ and $E ||Y_i||^2$ are bounded, then $\hat{\gamma}_n \xrightarrow{p} \gamma$ as $n \to \infty$. (ii) If additionally, Assumptions 4 and 7 hold, and if $E ||H_i||^2$, $E ||\beta_i||^2$, $E ||\widetilde{X}_i||^4$, and $E ||\widetilde{Z}_i||^4$ are bounded, then $\hat{\kappa} \xrightarrow{p} \kappa$ as $n \to \infty$.

 $^{^{9}}$ See Arellano and Bonhomme (2012); Chamberlain (1992); Graham and Powell (2012). See Laage (2024) for an approach without exogeneous regressors.

For consistency of CITE, we do not need to restrict the joint distribution of (Z, X, ϵ) beyond the existence of certain moments. The proof of consistency is standard and can be found in Appendix A. Because distribution theory for ITE and CITE is standard, we omit the remaining proofs.

4.2 Inconsistency of ITE

Theorem 1 establishes that the exogeneity restrictions in Assumptions 6 and 7 guarantee the consistency of CITE for κ and γ . We now show that the exogeneity restrictions are not sufficient for the consistency of ITE for κ .

Consider the special case that H is scalar, and that there is no Z, so the reduced form simplifies to:

$$\widetilde{Y}_{it} = \widetilde{X}_{it}H_i\kappa + \left(\widetilde{U}_{it} + \widetilde{X}_{it}\epsilon_i\right).$$

The ITE is simply a linear regression of \tilde{Y} on $\tilde{X}H$. Consistency of OLS requires that the error term is orthogonal to the regressors. Given Assumption 6,

$$E\left(\widetilde{X}_{it}H_i\widetilde{U}_{it}\right) = 0.$$

However, under the exogeneity Assumptions 6 and 7 we may have

$$E\left(\widetilde{X}_{it}^2 H_i \epsilon_i\right) \neq 0.$$

It is straightforward to construct data generating processes where this happens, see our simulation study in Section 6.

To ensure the consistency of ITE, one could further strengthen the conditions. For example, strengthening Assumption 7 to the following stronger *correlated random effects* assumption restores consistency:

Assumption 8 (Exogeneity, ϵ , strengthened). $E(\epsilon_i | H_i, X_i, H_i) = 0$.

This restriction requires the unobserved effect heterogeneity ϵ_i to be orthogonal to X and Z in addition to H. Consistency of ITE hence requires that unobserved effect heterogeneity is independent of all relevant regressors in the model.

Remark 2. Comparison of Assumptions 7 and 8 facilitates a key insight of our paper: the latter (which ensures consistency of ITE) is more demanding than Assumption 7 (which is sufficient for consistency of CITE), since it requires additional orthogonality of ϵ_i to X and Z (in addition to H). In the context of our empirical application, where we estimate how the effect of robots X on employment Y depends on interaction variables H, ITE requires that heterogeneity in the relationship between robotization and employment across countries is uncorrelated with countries' robot adoption. As discussed in our introduction, this is implausible because certain countries may be more open to new technologies, in which case they are more likely to adopt robots X and find ways to replace workers with robots, which gives rise to effect heterogeneity ϵ_i . CITE does not require such effect heterogeneity to be independent of robotization X for consistent estimation of the interaction effect.

5 Inference

In this section, we discuss how to conduct inference on the interaction effect κ for both estimators. For both estimators, inference is standard and can be performed using widely available software. We sketch our arguments below for the case where Z is absent. The general case with Z is similar, and it is omitted because the additional notation obscures the argument.

In this paper, we are concerned with inference for the partial effects. Empirical researchers may also be interested in the distribution of partial effects. Inference for this object is significantly more challenging, see for example Arellano and Bonhomme (2012) and Fernández-Val, Gao, Liao, and Vella (2022).

5.1 ITE

Inference for the ITE follows the standard theory for linear panel data models with clustered errors. In panel data, cluster-robust standard errors account for arbitrary correlation patterns in the error terms within each panel unit while maintaining the assumption of independence across units. Clustered standard errors address the fact that observations from the same panel unit (like multiple years of data from the same country) are typically not independent.

Two features of our model make clustering necessary. First, the within transformation induces correlation in the transformed errors \tilde{U}_i within each panel unit *i*. Second, the presence of unit-specific unobserved effect heterogeneity ϵ_i in the composite error term $V_i = \tilde{U}_i + \tilde{X}_i \epsilon_i$ creates an additional source of within-unit correlation. However, our random sampling assumption (Assumption 1) ensures independence across panel units, making cluster-robust standard errors appropriate.

The asymptotic variance of the ITE takes the standard sandwich form for clustered data:

$$\operatorname{var}(\check{\kappa}_n) = \left(\sum_i \Psi_i' \Psi_i\right)^{-1} \left(\sum_i \Psi_i' \operatorname{var}(V_i | \Psi_i) \Psi_i\right) \left(\sum_i \Psi_i' \Psi_i\right)^{-1}.$$

This variance can be consistently estimated in the usual way.

Implementation of cluster-robust standard errors is straightforward in standard statistical software. For instance, in current STATA versions, one would use the option vce(robust) for a standard panel command with interactions, like xtreg y x c.x#c.h, fe vce(robust). Similar commands are available in R, Python, and other commonly used statistical packages. Our simulation study in Section 6 documents the finite-sample performance of this inference approach.

5.2 CITE

Because of the two-step nature of CITE, it may be surprising that heteroskedasticityrobust standard errors in the second-step regression are sufficient for valid inference. To show that this is the case, recall that CITE for κ is a regression of the estimated individual-specific coefficients $\hat{\beta}_i$ on H_i ,

$$\widehat{\kappa}_n = \left(\sum_{i=1}^n H'_i H_i\right)^{-1} \sum_{i=1}^n H'_i \widehat{\beta}_i.$$

The corresponding regression equation is

$$\widehat{\beta}_i = \beta_i + (\widehat{\beta}_i - \beta_i) = H'_i \kappa + \epsilon_i + (\widehat{\beta}_i - \beta_i),$$

where the second equality follows from the heterogeneity equation (3).

Therefore, $\hat{\kappa}_n$ is based on a linear regression with a composite error term $\epsilon_i + (\hat{\beta}_i - \beta_i)$. This composite error term is independent across i, and heteroskedastic by construction. Independence follows from two observations: first, Assumption 1 implies that ϵ_i is independent across i, and second, the first-step estimation errors $\hat{\beta}_i = (\tilde{X}'_i \tilde{X}_i)^{-1} X'_i Y_i$ are independent across i because they only depend on $(\tilde{X}_i, \tilde{Y}_i)$, which are also independent across i by Assumption 1. The composite error term is heteroskedastic because the variance of the first-step estimation error $(\hat{\beta}_i - \beta_i)$ depends on \tilde{X}_i , even if ϵ_i is homoskedastic. Consequently, using heteroskedasticity-robust standard errors (White (1980)) in the second-step regression leads to valid inference.

Implementation is straightforward in standard statistical packages. For example, in STATA, after obtaining the unit-specific coefficients 'bhat' via a first-stage regression, one would use **reg bhat H**, **robust** to obtain valid standard errors. Our simulation study in Section 6 confirms good finite-sample performance of this approach. When control variables Z are present, the argument remains valid asymptotically, because γ is estimated at rate \sqrt{nT} . The contribution to sampling variation is negligible compared to that from $\hat{\beta}_i$. The latter is estimated using T observations.

6 Monte Carlo simulations

We use a Monte Carlo simulation to document the bias in ITE, the relative efficiency of ITE and CITE, and the finite sample performance of the inference procedures.

We use a data generating process in which all building blocks follow a normal distribution, are mutually independent, and are independent across *i*. For the heterogeneity equation, we let $H_i \sim \mathcal{N}(1,1)$, $\epsilon_i \sim \mathcal{N}(0,\sigma_{\epsilon})$, and $\beta_i = \kappa H_i + \epsilon_i$. For the outcome equation, we omit Z, and construct X_{it} as follows:

$$\psi_{it} \sim \mathcal{N}(0, \sigma_x), \quad \lambda_t \sim \mathcal{N}(0, \sigma_l),$$

 $X_{it} = (1 + \psi_{it})(1 + \Delta \times \epsilon_{it})(1 + \lambda_t).$

The parameter Δ controls the amount of correlation between X_{it} and ϵ_{it} . Values of Δ away from zero favor CITE because they create a correlation between X and unobserved effect heterogeneity ϵ (see Section 4.2). λ_t creates common time trends for all panel units. Finally, we generate $U_{it} \sim \mathcal{N}(0, \sigma_u)$, $\alpha_i \sim \mathcal{N}(0, \sigma_a)$, and compute $Y_{it} = \alpha_i + X_{it}\beta_i + U_{it}$.

All σ parameters default to 1. We will vary the number of time periods T and the number of cross-section units n, as well as the value of the interaction term coefficient κ and other design parameters.

6.1 The bias in ITE

We start with the design outlined above, with n = 100, T = 5, and $\kappa = 0.5$. Figure 1a plots the mean of ITE and CITE across 10000 simulations, as a function of $\Delta \in [-0.3, 0.3]$. At $\Delta = 0$, there is no endogeneity bias in the ITE, and both estimators

for κ have a mean of 0.5. For any $\Delta \neq 0$, the ITE is biased, whereas the mean CITE remains at 0.5. The bias in the ITE can be severe.

Figure 1b plots the mean of ITE and CITE as a function of $\kappa \in [-0.5, 0.5]$ at $\Delta = 0.4$. An unbiased estimator should follow the 45-degree diagonal, which is achieved by CITE. Conversely, the bias in ITE is substantial (≈ 0.3) and does not change much when varying κ . The shaded area corresponds to values of κ where the mean of ITE has the opposite sign of κ .



Figure 1: Mean of ITE and CITE.

6.2 Relative efficiency

Next, we compare the variability of ITE and CITE. We continue using n = 100, set $\Delta = 0$ (otherwise ITE is biased), and set $\kappa = 0$ (results for other values are comparable). We will vary $T \in \{3, \dots, 8\}$. In our baseliness results, the σ parameters are at 1. We will consider the effect of changing them on the relative efficiency of ITE and CITE as a function of T.

Figure 2a presents the baseline results. At T = 3, the ITE is more efficient than CITE. At T > 3, the CITE has a lower standard deviation. Figures 2b–2f document that the profiles shift as we vary the values of σ . For example, an increase in effect heterogeneity σ_{ϵ} favors the CITE. For example, at $\sigma_{\epsilon} = 0.1$ (Figure 2e), ITE is more efficient than CITE over the entire plotted range. Conversely, for $\sigma_{\epsilon} = 2$, CITE dominates ITE over T (Figure 2c).

Across all designs, the variability of both estimators decreases with T. The relative efficiency of the two estimators depends on the design, and there is no clear ranking that holds across all designs. Our main takeaway is that there is no clear penalty for using CITE to protect against bias.

6.3 Inference

In Section 5, we established simple inference procedures for ITE and CITE: use clusterrobust standard errors for ITE, and use heteroskedasticity-robust standard errors in the second step of CITE.

In both cases, the asymptotic distribution theory is standard, and standardized test statistics are asymptotically standard normal. Figure 3 plots the density estimates of the standardized test statistics across 100000 simulations in a design with n = 1000, T = 20, $\kappa = 0.5$, $\Delta = 0$, and all σ parameters equal to 1. The dashed line is the pdf of a standard normal distribution. The dashed line is barely visible in the plot, because the distribution of t-statistics of both ITE and CITE are right on top. For the same design, Figure 3b plots the density estimates of ITE and CITE. They appear bell-shaped, with CITE less variable than ITE. In conclusion, the results in Figure 3 are in line with the asymptotic theory.

Figure 4 provides further evidence on the asymptotic validity of our inferential procedure. We plot the rejection frequencies for H_0 : $\kappa_0 = 0.5$, varying κ in the data generating process along the horizontal axis. Both panels use 100000 simulations, and a nominal level of $\alpha = 0.05$ (indicated as a dashed horizontal line). Figure 4a presents results for a design with n = 100, T = 8, $\Delta = 0$, and all σ parameters equal to 1. In this design, ITE is consistent. At $\kappa = 0.5$, rejection frequencies are close to the nominal level of 0.05 for both estimators. Rejection frequences away from $\kappa = 0.5$ are better for CITE: an incorrect null hypothesis is more frequently rejected. Figure 4b presents results for a design with n = 500, T = 8, $\Delta = 0.3$, and all σ parameters equal to 1. In this design, the ITE curve is shifted left due to the bias in ITE under this design. Compared to the left panel, due to the increased sample size, the CITE curve is steeper around 0.5, and the rejection frequency at $\kappa = 0.5$ is very close to α .



Figure 2: The relative efficiency of CITE versus ITE changes depending on the design, and depends on T.



Figure 3: Validity of inference.



Figure 4: Rejection frequencies for ITE and CITE.

7 What does the relation between robots and employment depend on?

The relationship between robots and employment has attracted considerable attention over recent years. Jurat, Klump, and Schneider (2023); Mondolo (2022), and Restrepo (2024) survey this literature, which is highly relevant beyond academia: a solid knowledge about the labor market consequences of technical change helps us understand whose jobs may be automated, and to design targeted policies.

Yet, the empirical evidence on the robot-employment nexus is mixed to date. A meta-analysis by Guarascio, Piccirillo, and Reljic (2024) summarizes 33 studies with 644 estimates and finds considerable variation across studies and estimates, with a very small average negative relationship between robotization and employment. Within this literature, our application is related to cross-country industry level studies in the spirit of Graetz and Michaels (2018) and de Vries, Gentile, Miroudot, and Wacker (2020).

One possible reason for the mixed findings about the robot-employment relationship is cross-country heterogeneity, as suggested by previous empirical results by Reljic, Cirillo, and Guarascio (2023) and Chen and Frey (2024), for example. Such heterogeneity is highly plausible on theoretical grounds. Consider the seminal partial equilibrium model of Acemoglu and Restrepo (2020), where employment L in industry t and region i is affected by robotization R through three channels: a negative labor displacement effect of robotization technology, a positive demand effect reflecting the productivity gains from robotization, and a composition effect, which captures that industries that robotize require more (non-automated) labor from non-expanding sectors. Formally, those three channels are captured in their equation:

$$\frac{\Delta \ln L_{it}}{\Delta \ln R_{it}} = \underbrace{-\frac{\Delta \theta_t}{1 - \theta_t}}_{displacement} + \underbrace{\frac{1}{\alpha} \Delta \ln Y_i}_{demand} - \underbrace{(\sigma + \frac{1}{\alpha} - 1)\Delta \ln P_{it}^X}_{other},\tag{9}$$

where θ_t is a technology parameter indicating the range of tasks that can be automated (i.e., performed by robots), Y_c is output, $\sigma > 0$ is the elasticity of substitution (between goods of different industries) and P_{it} is the output price of industry t in i, and $1 - \alpha$ is the share of non-robot capital in the production process.¹⁰ Those effects may

¹⁰Acemoglu and Restrepo (2020) estimate a reduced-form version of equation (9) for US commuting zones i but their theoretical model is more general and can be applied across countries i. Note that our subscript t here indexes industries, not time periods. An earlier working paper version of our paper

operate into opposing directions and cause marginal effect heterogeneity. Heterogeneity across industries is thoroughly documented by Bekhtiar, Bittschi, and Sellner (2024), while our application focuses on heterogeneity across countries.

We focus on two key candidates for effect heterogeneity across countries in the robotemployment relationship, which arise from the partial equilibrium model by Acemoglu and Restrepo (2020): income p.c. levels and demand effects. The demand effect is straightforward and reflected in the second right-hand side term of equation (9). The importance of the income p.c. level for the robot-employment relationship arises from the fact that Acemoglu and Restrepo (2020) assume the technology parameter θ to be homogeneous across *i*, which is plausible for US commuting zones but not across countries. Moreover, θ captures whether a task can *technically* be automated. Whether this is *economically* profitable depends on the cost savings from using robots (π_i in Acemoglu & Restrepo, 2020, who for simplicity assume $\pi_i > 0 \forall i$). The cost-savings aspect unambiguously calls for a more negative effect of robotization on employment in higher-income countries: cost-saving potential π_i positively depends on wage levels, which are higher in high-income countries.

The remainder of this section hence explores how income p.c. and its changes, reflecting demand effects, alter the robot-employment relationship. We will abstract from other factors that may give rise the cross-country effect heterogeneity in the robot-employment nexus for the sake of providing a traceable and instructive illustration of the performance of CITE and ITE in a relevant applications with plausible interaction term effects. Code documentation is available on the authors' GitHub repository (https://github.com/KMWacker/CITErobots).

Regression setup: ITE vs. CITE

In cross-country industry-level studies, the relationship between labor market outcomes L and robots is typically estimated as a long-run first-difference equation (e.g., Bekhtiar et al., 2024; de Vries et al., 2020; Graetz & Michaels, 2018):

$$\Delta \ln L_{it} = \beta \Delta Robot s_{it} + a_i + \varepsilon_{it}.$$
(10)

The inclusion of country-specific fixed effects a_i is convenient for isolating country-

⁽Muris & Wacker, 2022) contains a textbook application from (Stock & Watson, 2015) that illustrates how ITE and CITE can differ in a standard panel setting with t indexing time periods.

specific employment trends that are associated with robots. To capture the effect heterogeneity suggested above, we interact $\Delta Robots$ with the relevant interaction terms H_1 (income levels) and H_2 (demand changes) in an ITE setting:

$$\Delta \ln L_{it} = \beta \Delta Robots_{it} + \sum_{k=1}^{K} \kappa_k \Delta Robots_{it} \cdot H_{k,i} + a_i + \nu_{it}, \qquad (11)$$

where our (up to) K = 2 interaction terms are country-specific income p.c. levels and demand changes. Recall that β in this framework corresponds to κ_0 in the context of CITE.

The standard ITE approach consists in estimating equation (11) through least squares (augmented with dummy variables α_i). As thoroughly discussed in Section 4, this requires strong exogeneity assumptions for consistent estimation of the interaction term coefficients κ_1, κ_2 . In particular, we may be concerned about two cases that give rise to a correlation between the error term u and the regressors.

One reason for concern is the behaviour of estimators in the presence of omitted interaction variables. Changes in demand H_{2i} are likely correlated with income levels H_{1i} and both may give rise to effect heterogeneity, leading to a bias in $\hat{\kappa}$ if one of them is omitted. We will hence assess to what extent it matters for ITE and CITE if $\kappa_1 H_1$ and $\kappa_2 H_2$ are jointly or individually included to gauge how both estimators behave in the presence of potentially omitted interaction variables.

A second reason for concern is more fundamental and concerns unobservable effect heterogeneity. For example, robots are frequently used for manual routine work in many industries, like health care. Patients, or clients more generally, in more technologyminded countries may be more open to being supported by robots. This gives rise to a higher worker substitutability, implying that the relationship between employment and robots is more negative in tech-minded countries. Tech-mindedness, which is not directly observable, will hence enter the error term u in our ITE regression equation (11) (through $X_{it}\epsilon_i$ in our general notation, see eq. (2)-(3)). Since tech-mindedness plausibly gives rise to higher robotization in the first place, the error term u and $\Delta Robots$ will be positively correlated in equation (11), which is a clear violation of the ITE exogeneity assumption (see Section 4.2).

Conversely, CITE explicitly models such country-specific effects. In the context of our application, CITE consists in first estimating the unit-specific parameters β_i for each individual country *i* through the outcome equation:

$$\Delta \ln L_{it} = \beta_i \Delta Robot s_{it} + a_i + u_{it}, \tag{12}$$

and subsequently exploring their relationship with country-specific income p.c. levels H_1 and demand changes H_2 by running the heterogeneity equation:

$$\widehat{\beta}_i = \kappa_0 + \kappa_1 H_{1i} + \kappa_2 H_{2i} + \epsilon_i.$$
(13)

Both approaches, ITE and CITE, allow us to investigate if previous empirical studies have omitted an essential factor how robots influence labor market outcomes and facilitate a comparison of the performances of ITE and CITE.

It is important to recall that CITE absorbs all country-specific effect heterogeneity in the robot-employment relation in β_i , while ITE only allows for heterogeneity in the modeled *interation* terms. Furthermore, note that equations (10) - (12) are longrun first-difference equations, which means that industry fixed effects are "differenced away" and that our subscript t now indexes industries, not time periods, to illustrate the applicability of CITE in various panel setups. A standard textbook panel application (as well as a higher-dimensional panel application) can be found in an earlier working paper version of our paper (Muris & Wacker, 2022).

Data

We use data from 15 manufacturing industries across 35 countries that is explained in full detail in Dijkstra and Wacker (2025). Robot stocks are calculated from IFR data (International Federabtion of Robotics) with a perpetual inventory method and divided by thousands of employees in the respective industry (from OECD TiM, see below). As suggested by Graetz and Michaels (2018), using raw or log changes in this robot density is not recommendable because of very low (or zero) starting values in the mid-2000s. We hence follow their recommendation to construct percentiles of (employment-weighted) changes in the robot stock.

For employment numbers, we rely on the 2023 edition of OECD's Trade in Employment (TiM) database (variable EMPN). Additionally, we merge the country-specific log of gross domestic product per capita, ln GDP p.c. (rgdpo/emp), from the Penn World Tables 10.0 (Feenstra, Inklaar, & Timmer, 2015), which captures income levels across countries and, in first differences, demand changes. Those are our interaction variables H_1, H_2 , as suggested by the first two right-hand side terms of equation (9). For $\Delta \ln L$ and $\Delta Robots$, we construct 10-year changes comparing the 2014-2018 averages to 2004-2008 averages.¹¹ For $\ln \text{GDP}$ p.c. we take country-specific averages over the 2004-2018 period. For changes in demand, we construct changes in $\ln \text{GDP}$ p.c. between the 2014-2018 and 2004-2008 averages (divided by 10 to obtain annual approximations). Further details on the data construction can be inferred from the code in the associated GitHub repository 'CITErobots'.

	Ν	mean	sd	\min	max
$\Delta \ln L$		-0.113	0.288	-1.39	0.92
Δ Robots		0.500	0.290	0.00	1.00
ln GDP p.c.		11.165	0.438	9.77	11.94
Δ demand		0.018	0.016	-0.01	0.06
Observations	509				

 Table 1: Summary Statistics

Figure 5: Changes in employment and robotization



Table 1 reports the summary statistics of our sample, while Figure 5 plots our key variables $\Delta \ln L$ and $\Delta Robots$ in a scatter that differentiates between lower- and

¹¹Comparing 10-year changes is standard in this literature. We construct 5-year averages at the start- and end-point to avoid year-specific variation in key variables (e.g., due to the global financial crisis in 2008).

higher-income countries. A number of features is worth highlighting. First, employment declines, on average (unweighted), with considerable heterogeneity: the standard deviation is more than double the mean. Second, since $\Delta Robots$ is constructed as percentile change, it ranges from 0 to 1 by construction, with a mean and IQR of 0.5. GDP p.c. is rather high, reflecting that most of our sample countries are high-income. Its rather low standard deviation reflects that this variable does not vary in the t dimension of our panel (i.e., across industries of the same country). Figure 5 separately illustrates countries at or above the ln GDP p.c. average (blue color) and those below average In GDP p.c. (black). Third, there is some negative correlation between $\Delta Robots$ and In GDP p.c., although this is not obviously visible from Figure 5. Their correlation coefficient is -0.10 (not reported) and averages of $\Delta Robots$ are slightly higher for lowerincome countries than for higher-income countries (0.54 vs. 0.47). Fourth, Figure 5 suggests a positive descriptive correlation between $\Delta \ln L$ and $\Delta Robots$ in lower-income countries (dashed black line) and no correlation in higher-income countries (solid blue line), which our estimation will explore more thoroughly. Fifth, the sample is slightly unbalanced (with 509 observations), reflecting missing values in 1-4 industries across 9 sample countries. Finally, it is worth noticing that changes in demand are close to the historic average GDP growth rate of 2% per annum documented in Pritchett and Summers (2014) and negatively correlated with ln GDP p.c. (-0.64, not reported), consistent with the idea of unconditional convergence (Patel, Sandefur, & Subramanian, 2021).

Note that our third and fourth observation, taken together and at face value, may constitute a problem for ITE if an interaction of $\Delta Robots$ with ln GDP p.c. is omitted: X ($\Delta Robots$) is plausibly correlated with the error term ϵ in the heterogeneity equation (3). This violates the required Assumption 8 for consistency of ITE and is likely on economic grounds: richer countries may exhibit certain features that simultaneously affect robotization and the robot-employment nexus.

Results

Table 2 summarizes our regression results, with ITE in the upper panel A and CITE in the lower panel B. We start with a simple linear regression of $\Delta \ln L$ on $\Delta Robots$ as a baseline in column (1), which facilitates comparison to the literature. In our (manufacturing) sample, there seems to be an overall positive relationship between industries' robotization and employment changes.

	(1)	(2)	(3)	(4)
	$\Delta \ln L$	$\Delta \ln L$	$\Delta \ln L$	$\Delta \ln L$
Panel A: ITE results				
Δ Robots	0.125^{***}	5.267^{***}	-0.008	4.053^{**}
Δ Robots \times ln GDP p.c.	(0.05)	(1.73) -0.454*** (0.15)	(0.00)	(1.97) -0.352^{**} (0.17)
Δ Robots \times Δ demand		(0.10)	10.475^{***} (3.77)	(0.17) 4.527 (5.03)
Country FEs	yes	yes	yes	yes
r2	0.308	0.326	0.320	0.327
Ν	509	509	509	509
Panel B: CITE results				
Δ Robots	0.449**	13.098*	0.214	16.380^{*}
	(0.18)	(7.33)	(0.16)	(9.36)
Δ Robots × ln GDP p.c.	· · · ·	-1.133*		-1.408*
		(0.65)		(0.81)
Δ Robots $\times \Delta$ demand			13.271	-11.892
			(9.60)	(14.99)
Country FEs	yes	yes	yes	yes
r2 (second stage)	N/A	0.209	0.037	0.226
N	35	35	35	35

Table 2: Regression results

* p<0.1, ** p<0.05, *** p<0.01

In column (2) of Table 2, we allow this relationship to vary with the interaction term for income levels (ln GDP p.c., H_1). In line with the above theoretical considerations (and already suggested by Figure 5), we find a negative interaction term coefficient $\hat{\kappa}_1$: the higher a country's income p.c. (and hence wage) level, the less favorable employment outcomes are associated with robotization. We will revert to the estimated marginal effects (which are graphically depicted in Figure 6) below but highlight a higher (absoulte) interaction term coefficient estimate for CITE, compared to ITE. Both estimators are relatively precise (interaction term standard errors < $1.7 \times |\hat{\kappa}_1|$), especially ITE.

In column (3) of Table 2, we explore the hypothesis that heterogeneities in the robot-

employment nexus are driven by demand changes H_2 . Consistent with theory both point estimates of the interaction term coefficient suggest a more positive association between robotization and employment growth the more aggregate demand increases in a country. However, ITE would leave the researcher highly confident in this specification that the partial effect between robots and employment linearly depends on demand effects. The associated ITE p-value < 0.01 in column (3) suggests that the data are quite compatible with the specified interaction model, while CITE inference raises considerably more caution. Also notice the low R squared of the second stage CITE projection in column (3).

Column (4), which includes both interaction terms, illustrates that the demand channel identified in column (3) is most likely due to an omitted variable bias. Recall that demand changes and income p.c. levels are inversely correlated across countries (-0.64): richer countries experience lower increases in demand. The interaction term $\Delta Robots \times \Delta demand$ in column (3) hence captures all cross-country heterogeniety that is correlated with demand changes, including substantial heterogeneity due to income p.c. levels. It is, of course, highly likely that $\Delta demand$ (and ln GDP p.c.) are correlated with other sources of effect heterogeneity in the robot-employment nexus.¹² As we know from our Monte Carlo inference simulations, ITE's rejection probabilities for t-tests of a null hypothesis about the interaction term coefficient can be highly distorted in this case. The stark ITE-related differences in p-values for such rejections when moving from column (3) to column (4) in panel A of Table 2 should be a warning signal for applied researchers that inference of standard ITE may be highly susceptible and misleading.

Marginal effects

The results in Table 2 suggest that cross-country heterogeneity in the robot-employment relationship is to a significant degree driven by differences in income per capita (and associated wage) levels. Column (2) of panel B suggests that 21% of cross-country heterogeneity can be explained by variation of ln GDP p.c. (see also Figure 6).

How much of a difference does the income level make for employment effects of robotization?¹³ And how would that estimated marginal effect vary between ITE and CITE? Figure 6 is instructive for answering this question. Let us consider income levels

¹²Likewise, it is plausible that our measure for demand changes is erroneous – which substantiates our point that ITE inference in column (3) of panel A in Table 2 is problematic.

¹³In this section, we treat the estimated relationship as an 'effect', well-aware that other sources of endogeneity may pose challenges to a causal interpretation.



Figure 6: Marginal effects of robots on employment

of $\ln \text{GDP p.c.} = (10.5, 11.4, 11.7)$. In our sample, the former corresponds to Bulgaria, which is at the lower end of new European Union member states and a plausible level for various middle-upper income countries (e.g., between Colombia and Chile). The second corresponds to Germany, the latter to the United States. Recall that an overall regression with no interaction effects would suggest a positive relationship equal to 0.13 in either case (red line in Figure 6 and column (1) in panel A of Table 2).

Considering a difference in robotization rates of $\Delta Robots = 0.5$ and estimates from column (2) of Table 2, ITE suggests an associated employment increase of 25% over a decade, for a country at a ln GDP level of 10.5.¹⁴ The magnitude is much larger with CITE, 60%, and Figure 6 suggests that it more plausibly captures the estimated country-specific coefficients for various upper-middle income countries, while ITE seems heavily impacted by the low country-specific estimate of China, the lowest-income country in our sample.

For a ln GDP p.c. level of 11.4 (approximately Germany), the estimated marginal effects of CITE (9.1%) and ITE (4.6%) are both positive and relatively close to the

¹⁴Recall that $\Delta Robots$ is constructed as percentile changes, ranging from 0 to 1. The IQR of $\Delta Robots = 0.5$ hence compares the 25th percentile of robot adoption to the 75th percentile of robot adoption. Estimated marginal effect: $0.5 \cdot (\hat{\kappa}_0 + \hat{\kappa}_1 \ln GDPp.c.)$.

magnitude without interaction effect ($6\% = 0.5 \cdot 0.13$, cf. column (1) panel A of Table 2). For a ln GDP p.c. level of 11.7 (approximately US), the estimated marginal effects are negative, and stronger for CITE (-7.9%) than for ITE (-2.2%). Considering that CITE is consistent under milder assumptions than ITE, our results suggest that ITE and homogeneous estimation without interaction terms both underestimate the negative effects of robotization for manufacturing employment in countries at the highest income levels in the sample (ln GDP p.c. > 11.57; Ireland, Norway, Switzerland, United States).

If income (and associated wage) levels are indeed a main source of robot-employment effect heterogeneity, this may reconcile conflicting country level evidence in the literature. For example, Acemoglu and Restrepo (2020), find negative consequences of robot exposure on regional labor markets in the US, while Dauth, Findeisen, Suedekum, and Woessner (2021) find no such aggregate employment effects for Germany, despite using a similar approach. This is quite consistent with our quantitative estimates about effect heterogeneity due to income differences across countries (especially when considering differences in sample coverage and methodologies).

CITE, of course, provides country-specific estimates of the robot-employment relationship, as they are illustrated in Figure 6. The reason why we do not consider those country-specific estimates for our marginal effect calculations is that they suffer from incidental parameter bias. The purpose of our paper is to estimate interaction terms in exactly such a short-T panel, where country-specific estimation is no option, and we have shown that CITE offers a consistent projection of those (individually biased) panel unit estimates on the interaction variable of interest. We hence consider marginal effects of a random country at an income level like, e.g., Bulgaria, rather than a Bulgaria-specific coefficient. Applied analysts working on specific countries may be willing to accept the incidental parameter bias in country-specific estimates, but this is not the purpose of our paper.

Application takeaways

An important insight of our analysis for the literature on employment effects of robotization is that economic displacement effects of robots are plausibly stronger in countries at higher income levels, where wages are usually higher. Econometrically, our application suggests that this income-interaction effect is underestimated by standard ITE and that ITE may lead to premature rejection of a null hypothesis about an interaction term in the presence of misspecification (e.g., when another relevant interaction term is omitted).

Our application results, and particularly Figure 6, also highlight why CITE is attractive for applied econometricians. Because unit-specific coefficients are estimated in a first step and then further analyzed in a second step (through projection on the relevant interaction variables H), this offers a lot of flexibility and diagnostics tools in the second step.¹⁵

8 Conclusion

Our paper has highlighted that the most common approach to estimate interaction effects with panel data neglects residual heterogeneity in the effect of X on Y across panel units. If this unobserved effect heterogeneity is systematically correlated with explanatory variables in the model, which is plausible in most econometric applications, this standard approach will not provide consistent estimates of the interaction effects.

We advocate for explicit modeling of effect heterogeneity across panel units with a correlated interaction term estimator (CITE) that requires less demanding exogeneity assumptions for consistency. Those findings are irrespective of whether interaction variables are constant within panels or not.

Our argument in favor of CITE resembles the case for additive fixed effects in linear panel data models. While additive fixed effects explicitly capture unobserved heterogeneity in unit-specific levels of Y, CITE explicitly models the effect heterogeneity β_i that most common panel interaction estimates do not account for.

Our paper has focused on the comparison of ITE and CITE for the estimation of interaction effects in panel data. Two alternative approaches that are sometimes used in empirical practice are (i) subgroup analysis – in which one conducts the analysis conditional on every value of the interaction variable (e.g. Chetty, Hendren, and Katz (2016)) – and (ii) saturating the heterogeneity model and then running ITE (e.g., Norris, Pecenco, and Weaver (2021)). Either approach may bridge the gap between ITE and CITE. Although such approaches may work well when the interaction variables are discrete, they are not feasible or practical for continuous interaction variables. It would

¹⁵For example, one may speculate to what extent Bulgaria, Colombia, China, and Chile are outliers. Indeed, those countries surpass a critical value for 'Cook's distance' and when they are excluded from the analysis, we obtain ITE and CITE interaction coefficients for ln GDP p.c. of -0.58 and -0.59 (significant at the 1 and 10% level), respectively, which are very close to each other and fall between the point estimates reported in column (2) of Table 6.

be interesting to extend our analysis to include a comparison with these alternative approaches.

An obvious extension of our research is to derive a test statistic for consistency of ITE when compared to CITE. Until such a test becomes available, we strongly recommend the use of CITE or at least reporting CITE results next to ITE results. Given the simple implementation of CITE, this should become standard in applied econometrics when estimating interaction terms in panel data.¹⁶

An additional advantage of CITE, beyond its less restrictive exogeneity assumptions, is that the unit-specific effect heterogeneity captured in β_i can provide further insights about sources of heterogeneity in panel data. Recent advances in machine learning appear particularly promising to further explore those first-step CITE estimates of unit-specific effects.

¹⁶If one is interested in interactions with varying variables G, it is sufficient to additionally include an interaction of the X variable(s) with panel dummy variables. If one is interested in interactions with H variables that are constant within panels, those panel specific estimates from the first step need to be saved and regressed on the H variables in a second step. An adjustment of standard errors is necessary to account for first-step estimation.

References

- Acemoglu, D., & Restrepo, P. (2020). Robots and Jobs: Evidence from US Labor Markets. Journal of Political Economy, 128(6), 2188-2244.
- Alfaro, L., Chanda, A., Kalemli-Ozcan, S., & Sayek, S. (2004). FDI and economic growth: The role of local financial markets. *Journal of International Economics*, 64(1), 89-112.
- Alfaro, L., Chanda, A., Kalemli-Ozcan, S., & Sayek, S. (2010). Does foreign direct investment promote growth? Exploring the role of financial markets on linkages. *Journal of Development Economics*, 91(2), 242–256.
- Alsan, M., & Goldin, C. (2019). Watersheds in child mortality: The role of effective water and sewerage infrastructure, 1880–1920. Journal of Political Economy, 127(2), 586–638.
- Amiti, M., & Konings, J. (2007). Trade liberalization, intermediate inputs, and productivity: Evidence from indonesia. American Economic Review, 97(5), 1611–1638.
- Arellano, M., & Bonhomme, S. (2012). Identifying Distributional Characteristics in Random Coefficients Panel Data Models. *Review of Economic Studies*, 79(3), 987–1020.
- Balli, H. O., & Sørensen, B. E. (2013). Interaction effects in econometrics. *Empirical Economics*, 45(1), 583–603.
- Bekhtiar, K., Bittschi, B., & Sellner, R. (2024). Robots at work? Pitfalls of industrylevel data. *Journal of Applied Econometrics*, 39(6), 1180-1189.
- Berman, N., Martin, P., & Mayer, T. (2012). How do different exporters react to exchange rate changes? Quarterly Journal of Economics, 127(1), 437–492.
- Bloom, N., Draca, M., & Van Reenen, J. (2016). Trade induced technical change? The impact of chinese imports on innovation, IT and productivity. *Review of Economic Studies*, 83(1), 87–117.
- Bloom, N., Sadun, R., & Van Reenen, J. (2012). Americans do IT better: US multinationals and the productivity miracle. American Economic Review, 102(1), 167–201. doi: 10/gftrv4
- Borensztein, E., Gregorio, J. D., & Lee, J.-W. (1998). How does foreign direct investment affect economic growth? Journal of International Economics, 4(1), 115 -135.
- Burnside, C., & Dollar, D. (2000). Aid, Policies, and Growth. American Economic Review, 90(4), 23.

- Chamberlain, G. (1992). Efficiency Bounds for Semiparametric Regression. *Economet*rica, 60(3), 567–596.
- Chen, C., & Frey, C. B. (2024). Robots and reshoring: a comparative study of automation, trade, and employment in Europe. *Industrial and Corporate Change*, 33(6), 1331-1377.
- Chetty, R., Hendren, N., & Katz, L. F. (2016). The Effects of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportunity Experiment. American Economic Review, 106(4), 855–902.
- Couttenier, M., Petrencu, V., Rohner, D., & Thoenig, M. (2019, December). The Violent Legacy of Conflict: Evidence on Asylum Seekers, Crime, and Public Policy in Switzerland. *American Economic Review*, 109(12), 4378–4425.
- Dauth, W., Findeisen, S., Suedekum, J., & Woessner, N. (2021). The Adjustment of Labor Markets to Robots. *Journal of the European Economic Association*, 19(6), 3104-3153.
- de Chaisemartin, C., & D'Haultfœuille, X. (2022). Two-Way Fixed Effects and Differences-in-Differences with Heterogeneous Treatment Effects: A Survey (Working Paper No. 29691). National Bureau of Economic Research.
- de Chaisemartin, C., & D'Haultfœuille, X. (2024). Credible Answers to Hard Questions: Differences-in-Differences for Natural Experiments.
- de Vries, G. J., Gentile, E., Miroudot, S., & Wacker, K. M. (2020). The rise of robots and the fall of routine jobs. *Labour Economics*, 66, 101885.
- de Chaisemartin, C., & Lei, Z. (2023). More Robust Estimators for Instrumental-Variable Panel Designs, With An Application to the Effect of Imports from China on US Employment (No. arXiv:2103.06437). arXiv. doi: 10.48550/arXiv.2103 .06437
- Dijkstra, H., & Wacker, K. M. (2025). Robots, shoring patterns, and employment [Working Paper]. online.
- Duflo, E., Dupas, P., & Kremer, M. (2011). Peer Effects, Teacher Incentives, and the Impact of Tracking: Evidence from a Randomized Evaluation in Kenya. American Economic Review, 101(5), 1739–1774.
- Durham, J. B. (2004). Absorptive capacity and the effects of foreign direct investment and equity foreign portfolio investment on economic growth. *European Economic Review*, 48(2), 285-306.
- Epifani, P., & Gancia, G. (2009). Openness, government size and the terms of trade. *Review of Economic Studies*, 76(2), 629–668.

- Feenstra, R. C., Inklaar, R., & Timmer, M. P. (2015). The next generation of the penn world table. American Economic Review, 105(10), 3150-82.
- Fernández-Val, I., Gao, W. Y., Liao, Y., & Vella, F. (2022). Dynamic heterogeneous distribution regression panel models, with an application to labor income processes. arXiv preprint arXiv:2202.04154.
- Fernández-Val, I., & Lee, J. (2013). Panel data models with nonadditive unobserved heterogeneity: Estimation and inference. *Quantitative Economics*, 4(3), 453-481.
- Giesselmann, M., & Schmidt-Catran, A. W. (2020). Interactions in Fixed Effects Regression Models. Sociological Methods & Research, 51(3).
- Graetz, G., & Michaels, G. (2018, 12). Robots at Work. *Review of Economics and Statistics*, 100(5), 753-768.
- Graham, B. S., & Powell, J. L. (2012). Identification and Estimation of Average Partial Effects in "Irregular" Correlated Random Coefficient Panel Data Models. *Econometrica*, 80(5), 2105–2152.
- Guarascio, D., Piccirillo, A., & Reljic, J. (2024). Will robot replace workers? Assessing the impact of robots on employment and wages with meta-analysis (Tech. Rep.).
- Herrera, H., Ordoñez, G., & Trebesch, C. (2020). Political booms, financial crises. Journal of Political Economy, 128(2), 507–543.
- Jurat, A., Klump, R., & Schneider, F. (2023). *Robots and wages: A meta-analysis* (Tech. Rep.). ZBW Leibniz Information Centre for Economics, Kiel, Hamburg.
- Laage, L. (2024). A Correlated Random Coefficient Panel Model with Time-Varying Endogeneity. *Journal of Econometrics*, 242(2), 105804.
- List, J. A., & Sturm, D. M. (2006). How elections matter: Theory and evidence from environmental policy. *Quarterly Journal of Economics*, 121(4), 1249–1281.
- MaCurdy, T. E. (1981). An Empirical Model of Labor Supply in a Life-Cycle Setting. Journal of Political Economy, 89(6), 1059–1085.
- Manacorda, M., & Tesei, A. (2020). Liberation Technology: Mobile Phones and Political Mobilization in Africa. *Econometrica*, 88(2), 533–567.
- Mondolo, J. (2022). The composite link between technological change and employment: A survey of the literature. *Journal of Economic Surveys*, 36(4), 1027-1068.
- Muris, C., & Wacker, K. M. (2022). Estimating interaction effects with panel data (Tech. Rep. No. arXiv:2211.01557 [econ.EM]).
- Murtazashvili, I., & Wooldridge, J. M. (2008). Fixed effects instrumental variables estimation in correlated random coefficient panel data models. *Journal of Econometrics*, 142(1), 539–552.

- Norris, S., Pecenco, M., & Weaver, J. (2021). The Effects of Parental and Sibling Incarceration: Evidence from Ohio. American Economic Review, 111(9), 2926– 2963.
- Patel, D., Sandefur, J., & Subramanian, A. (2021). The new era of unconditional convergence. Journal of Development Economics, 152, 102687.
- Pritchett, L., & Summers, L. H. (2014). Asiaphoria meets regression to the mean (NBER Working Paper No. 20573). National Bureau of Economic Research.
- Reljic, J., Cirillo, V., & Guarascio, D. (2023). Regimes of robotization in Europe. Economics Letters, 232, 111320.
- Restrepo, P. (2024). Automation: Theory, evidence, and outlook. Annual Review of Economics, 16, 1-25.
- Sasaki, Y., & Ura, T. (2021). *Slow Movers in Panel Data* (No. arXiv:2110.12041). arXiv:2110.12041. arXiv.
- Shambaugh, J. C. (2004). The effect of fixed exchange rates on monetary policy. Quarterly Journal of Economics, 119(1), 301-352.
- Słoczyński, T. (2022). Interpreting OLS Estimands When Treatment Effects Are Heterogeneous: Smaller Groups Get Larger Weights. The Review of Economics and Statistics, 104(3), 501–509.
- Spilimbergo, A. (2009). Democracy and Foreign Education. American Economic Review, 99(1), 528–543.
- Stock, J. H., & Watson, M. W. (2015). Introduction to econometrics (Updated third edition, global edition ed.). Boston Columbus Indianapolis New York San Francisco Hoboken Amsterdam Cape Town Dubai London: Pearson.
- Storeygard, A. (2016). Farther on down the Road: Transport Costs, Trade and Urban Growth in Sub-Saharan Africa. *Review of Economic Studies*, 83(3), 1263–1295.
- Verdier, V. (2020). Average treatment effects for stayers with correlated random coefficient models of panel data. Journal of Applied Econometrics, 35(7), 917– 939.
- White, H. (1980). A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity. *Econometrica*, 48(4), 817.
- Winkelmann, R. (2024). Neglected Heterogeneity, Simpson's Paradox, and the Anatomy of Least Squares. Journal of Econometric Methods, 13(1), 131–144.
- Wooldridge, J. M. (2005, May). Fixed-Effects and Related Estimators for Correlated Random-Coefficient and Treatment-Effect Panel Data Models. The Review of Economics and Statistics, 87(2), 385–390. doi: 10.1162/0034653053970320

A Proof of Theorem 1

Proof. Part (i): Consistency of $\hat{\gamma}_n$. CITE for γ is the coefficient estimate in a linear regression of $M_i \tilde{Y}_i$ on $M_i \tilde{Z}_i$, corresponding to

$$M_i \widetilde{Y}_i = M_i \widetilde{Z}_i \gamma + M_i \widetilde{U}_i$$

It follows that

$$\widehat{\gamma}_n - \gamma = \left(\frac{\sum_{i=1}^n \widetilde{Z}'_i M_i \widetilde{Z}_i}{n}\right)^{-1} \frac{\sum_{i=1}^n \widetilde{Z}'_i M_i \widetilde{U}_i}{n}.$$

The boundedness assumptions in the statement of the theorem, and the fact that the residual maker matrices M_i and the linear transformation $Y_i \to \widetilde{Y}_i$ have bounded norm, imply that $E[\widetilde{Z}'_i M_i \widetilde{Z}_i]$ and $E[\widetilde{Z}'_i M_i \widetilde{U}_i]$ are bounded. Thus, the weak law of large numbers (WLLN) for random vectors ensures that both terms converge to their expectations. Because $E[\widetilde{Z}'_i M_i \widetilde{Z}_i]$ is invertible by Assumption 3, it only remains to be shown that $E\left(\widetilde{Z}'_i M_i \widetilde{U}_i\right) = 0$. Assumption 6 implies $E\left(\widetilde{U}_i \middle| \widetilde{X}_i, \widetilde{Z}_i\right) = 0$. By the law of iterated expectations, $E\left(\widetilde{Z}'_i M_i \widetilde{U}_i\right) = 0$, since M_i is a transformation of \widetilde{X}_i .

Part (ii): Consistency of $\hat{\kappa}_n$.

The CITE for κ , cf. Definition 2, has the representation

$$\widehat{\kappa}_n = \left(\sum_{i=1}^n H'_i H_i\right)^{-1} \sum_{i=1}^n H'_i \widehat{\beta}_i$$
$$= \left(\sum_{i=1}^n H'_i H_i\right)^{-1} \sum_{i=1}^n H'_i \beta_i + \left(\sum_{i=1}^n H'_i H_i\right)^{-1} \sum_{i=1}^n H'_i \left(\widehat{\beta}_i - \beta_i\right).$$

Because the second moments of H and β are assumed bounded in the statement of the theorem, and Assumption 4 guarantees invertibility of E[H'H], the WLLN applied to the first term yields κ by Assumption 7.

If we can show that the second term converges to zero in probability, we have established consistency of $\hat{\kappa}_n$ to κ . Recall the individual specific effects estimators

$$\hat{\beta}_{i} = \left(\widetilde{X}_{i}'\widetilde{X}_{i}\right)^{-1}\widetilde{X}_{i}'\left(\widetilde{Y}_{i} - \widetilde{Z}_{i}\widehat{\gamma}_{n}\right)$$
$$= \left(\widetilde{X}_{i}'\widetilde{X}_{i}\right)^{-1}\widetilde{X}_{i}'\left(\widetilde{X}_{i}\beta_{i} + \widetilde{Z}_{i}\left(\gamma - \widehat{\gamma}_{n}\right) + \widetilde{U}_{i}\right),$$

where the last step uses (6). Thus, we can write

$$\widehat{\beta}_i - \beta_i = \left(\widetilde{X}'_i \widetilde{X}_i\right)^{-1} \widetilde{X}'_i \widetilde{Z}_i \left(\gamma - \widehat{\gamma}_n\right) + \left(\widetilde{X}'_i \widetilde{X}_i\right)^{-1} \widetilde{X}'_i \widetilde{U}_i.$$

This allows us to write

$$\begin{aligned} \widehat{\kappa}_n - \kappa &= \left(\sum_{i=1}^n H'_i H_i\right)^{-1} \sum_{i=1}^n H'_i \left(\widehat{\beta}_i - \beta_i\right) \\ &= \left(\sum_{i=1}^n H'_i H_i\right)^{-1} \left[\sum_{i=1}^n H'_i (\widetilde{X}'_i \widetilde{X}'_i)^{-1} \widetilde{X}'_i \widetilde{Z}_i\right] (\gamma - \widehat{\gamma}_n) + \\ &+ \left(\sum_{i=1}^n H'_i H_i\right)^{-1} \sum_{i=1}^n H'_i (\widetilde{X}'_i \widetilde{X}'_i)^{-1} \widetilde{X}'_i \widetilde{U}_i \\ &\equiv A_{1n} + A_{2n}. \end{aligned}$$

Assumption 2 ensures that $(\widetilde{X}'_i \widetilde{X}'_i)^{-1} \leq \frac{1}{h^2}$ uniformly over *i*, so the term can be ignored in the analysis of A_{1n} and A_{2n} . Applying Cauchy-Schwartz twice to each term, and using: boundedness of the second moments of *H*; boundedness of the fourth moments of \widetilde{X} and \widetilde{Z} ; that $\widehat{\gamma}_n \xrightarrow{p} \gamma$ from part (i); and that $E[H'_i \widetilde{X}'_i \widetilde{U}_i]$ from Assumption 6, we conclude that A_{1n} and A_{2n} converge in probability to zero.