

Low-rank approximations of nonseparable panel models

IVÁN FERNÁNDEZ-VAL[†], HUGO FREEMAN[‡] AND MARTIN WEIDNER[§]

[†]*Department of Economics, Boston University, 270 Bay State Road, Boston, MA 02215-1403, USA.*

Email: ivanf@bu.edu

[‡]*Department of Economics, University College London, Gower Street, London WC1E 6BT, UK.*

Email: hugo.freeman.16@ucl.ac.uk

[§]*Nuffield College and Department of Economics, University of Oxford, Manor Road, Oxford OX1 3UQ, UK.*

Email: martin.weidner@economics.ox.ac.uk

First version received: 22 October 2020; final version accepted: 4 March 2021.

Summary: We provide estimation methods for nonseparable panel models based on low-rank factor structure approximations. The factor structures are estimated by matrix-completion methods to deal with the computational challenges of principal component analysis in the presence of missing data. We show that the resulting estimators are consistent in large panels, but suffer from approximation and shrinkage biases. We correct these biases using matching and difference-in-differences approaches. Numerical examples and an empirical application to the effect of election day registration on voter turnout in the US illustrate the properties and usefulness of our methods.

Keywords: *Nonseparable panel, low-rank approximations, matrix completion, debias, two-way matching, election day registration.*

JEL codes: *C14, C23.*

1. INTRODUCTION

Nonseparable models are useful to capture multidimensional unobserved heterogeneity, which is an important feature of economic data. The presence of this heterogeneity makes the effect of covariates on the outcome of interest different for each unit due to factors that are unobservable or unavailable to the researcher. In the absence of further restrictions, a different data generating process essentially operates for each unit, which creates identification and estimation challenges. One way to deal with these challenges is the use of panel data, where each unit is observed on multiple occasions. In this paper, we develop an approach to estimate nonseparable models from panel data based on homogeneity restrictions and low-rank factor approximations. While homogeneity restrictions have been used previously in this context, the application of low-rank factor approximations is more novel.

The nonseparable model that we consider includes observed discrete covariates or treatments, multidimensional unobserved individual and time effects, and idiosyncratic errors. We construct the effects of interest as averages or quantiles of potential outcomes constructed from the model by exogenously manipulating the value of the treatments. These effects are generally not identified

from the observed data because the treatment assignment is usually determined by the unobserved individual and time effects. Following the previous panel literature, we impose cross-section and time-series homogeneity restrictions to identify the effects of interest; see, e.g., Chamberlain (1982), Manski (1987), Honoré (1992), Evdokimov (2010), Graham and Powell (2012), Hoderlein and White (2012) and Chernozhukov et al. (2013).

The estimation of the nonseparable model is challenging due to the presence of the multidimensional unobserved individual and time effects. We cannot just exclude these effects because they are endogenous, i.e., related to the treatments. We deal with this problem by approximating their effect with a low-rank factor structure. This approach can be interpreted as a series or sieve approximation on the unobservables. We characterize the error of this approximation in terms of the functional singular value decomposition of the expectation of the outcome conditional on the treatment and unobserved effects. For smooth conditional expectation functions, the mean squared error of the approximation error vanishes with the rank of the factor structure at a polynomial rate.

We develop an estimator of the low-rank factor approximation in the case where the covariate of interest is binary. This is an empirically relevant case as it covers the treatment effect model for panel data. We also show how to extend the model to include additive controls and fixed effects. Here, we rely on the analogy between the estimation of treatment effects and the matrix-completion problem previously noted by Athey et al. (2021) and Amjad et al. (2018). Thus, given that the principal components program is combinatorially hard in the presence of missing data, we consider the convex relaxation of this program that replaces a constraint in the rank of a matrix by a constraint in its nuclear norm, following Fazel (2003) and Srebro and Jaakkola (2003). The resulting estimator is the matrix-completion estimator.

The main theoretical result of the paper is to show that the matrix-completion estimator is consistent under asymptotic sequences where the two dimensions of the panel grow to infinity at the same rate. This result does not follow from the existing matrix completion literature that assumes that the matrix to complete has low rank. In our case, the underlying matrix of interest can have full rank, but we impose appropriate smoothness assumptions on the data generating process that guarantee that the singular values of the matrix form a rapidly decreasing sequence. This allows a low-rank approximation, and it also implies a bound on the nuclear norm of the matrix. Our consistency proof for the matrix-completion estimator therefore crucially relies on the bound of the nuclear norm, but does not impose any low-rank conditions. Our proof strategy also avoids the high-level *restricted strong convexity* assumption (see, e.g., Negahban and Wainwright, 2012). We instead provide interpretable conditions on the underlying process of the observable and unobservable variables directly.

The matrix-completion estimator is consistent, but can be biased in small samples. This bias comes from two different sources: approximation bias due to the low-rank factor structure approximation; and shrinkage bias due to the nuclear norm regularization of the principal component analysis program (Cai et al., 2010; Ma et al., 2011; Bai and Ng, 2019b). We propose matching approaches to debias the estimator. For each treatment level, the simplest approach consists of finding the observation in the other treatment level that is the closest in terms of the estimated factor structure. We also propose a two-way matching procedure that combines matching with a differences-in-differences approach. The two-way procedure is related to several recent proposals, such as the matching approach of Imai and Kim (2019) to estimate causal effects from panel data and the blind regression of Li et al. (2017) for matrix completion. The difference with these proposals is in the information used to match the observations. Imai and Kim (2019) use the treatment variable and Li et al. (2017) the outcome, whereas we use the estimated factor

structure. In this sense, the estimation of the factor structure can be seen as a preliminary de-noising step of the data as in Chatterjee (2015). Amjad et al. (2018) proposed a similar debiasing procedure based on the estimated factor structure, but they rely on synthetic control methods instead of matching. In contemporaneous and independent work, Chernozhukov et al. (2020) have developed an alternative rotation-debiasing method that can be applied to make inference on heterogeneous treatment effects in low-rank models. This method consists of the application of iterative least squares to the left and right singular vectors of the matrix-completion estimator.

We illustrate our methods with an empirical application to the effect of election day registration (EDR) on voter turnout and numerical simulations. We estimate average and quantile effects using a state-level panel dataset on the 24 US presidential elections between 1920 and 2012 collected by Xu (2017). We find that, after controlling for possible nonrandom adoption, EDR has a positive effect, especially at the bottom of the voter turnout distribution. Our methods uncover stronger effects than standard difference-in-differences methods that rely on restrictive parallel trend assumptions. The simulation results show that our theoretical results provide a good representation of the behaviour of the estimators in small samples.

The rest of the paper is organized as follows. Section 2 describes the model and effects of interest. Section 3 introduces the low-rank factor approximation and derives the properties of its matrix-completion estimator. The matching methods to debias the matrix-completion estimator are discussed in Section 4. Section 5 reports the results of the numerical examples. All the proofs of the theoretical results are gathered in the Appendix.

2. MODEL AND EFFECTS OF INTEREST

Throughout this paper we consider the following nonseparable and nonparametric panel data model:

ASSUMPTION 2.1. (MODEL)

$$Y_{it} = g(\mathbf{X}_{it}, \mathbf{A}_i, \mathbf{B}_t, \mathbf{U}_{it}), \quad i \in \mathbb{N} = \{1, \dots, N\}, t \in \mathbb{T} = \{1, \dots, T\}, \quad (2.1)$$

where i and t index individual units and time periods, respectively; Y_{it} is an observed outcome or response variable with support $\mathbb{Y} \subseteq \mathbb{R}$; g is an unknown function; \mathbf{X}_{it} is a vector of observed covariates or treatments with finite support \mathbb{X} ; \mathbf{A}_i and \mathbf{B}_t are vectors of individual and time unobserved effects, possibly correlated with \mathbf{X}_{it} , with supports $\mathbb{A} \subseteq \mathbb{R}^{d_a}$ and $\mathbb{B} \subseteq \mathbb{R}^{d_b}$, respectively; and \mathbf{U}_{it} is a vector of unobserved error terms of unspecified dimension, for which we assume that

$$\mathbf{U}_{it} \stackrel{d}{=} \mathbf{U}_{js} \mid \mathbf{X}^{NT}, \mathbf{A}^N, \mathbf{B}^T, \quad \text{for all } i, j \in \mathbb{N}, t, s \in \mathbb{T}, \quad (2.2)$$

and

$$\mathbf{U}_{it} \perp\!\!\!\perp (\mathbf{X}^{NT}, \mathbf{A}^N, \mathbf{B}^T) \mid \mathbf{A}_i, \mathbf{B}_t, \quad \text{for all } i \in \mathbb{N}, t \in \mathbb{T}, \quad (2.3)$$

where $\mathbf{X}^{NT} = \{\mathbf{X}_{it} : i \in \mathbb{N}, t \in \mathbb{T}\}$, $\mathbf{A}^N = \{\mathbf{A}_i : i \in \mathbb{N}\}$, $\mathbf{B}^T = \{\mathbf{B}_t : t \in \mathbb{T}\}$, and $\perp\!\!\!\perp$ denotes stochastic independence. We also assume that, for all $i \in \mathbb{N}$, $t \in \mathbb{T}$, the support of $(\mathbf{X}_{it}, \mathbf{A}_i, \mathbf{B}_t)$ is equal to the Cartesian product $\mathbb{X} \times \mathbb{A} \times \mathbb{B}$, and that $E Y_{it}^2 < \infty$.

This model can be motivated from a purely statistical perspective as a latent variable model using the Aldous-Hoover representation for exchangeable random matrices, e.g., Xu et al. (2014),

Chatterjee (2015), Orbanz and Roy (2015), and Li and Bell (2017).¹ We motivate it instead as a structural model where the unobserved effects \mathbf{A}_i and \mathbf{B}_t are associated with individual heterogeneity and aggregate shocks, respectively. Additional exogenous covariates can be incorporated in the usual way by carrying out the analysis conditional on them. We focus on discrete covariates but, from a theoretical perspective, the extension to continuous covariates is straightforward by using appropriate smoothing methods—it is, however, not clear to us whether that extension would be practically useful with realistic sample sizes. We therefore think that it would complicate our presentation without much benefit.

The main restriction imposed by Assumption 2.1 is the unit and time homogeneity in (2.2). A sufficient condition for unit homogeneity is that the observations are identically distributed across i , which is a common sampling assumption for panel data. Time homogeneity has also been commonly used in panel data models (Chamberlain, 1982; Manski, 1987; Honoré, 1992; Evdokimov, 2010; Graham and Powell, 2012; Hoderlein and White, 2012; Chernozhukov et al., 2013). It implies that time is randomly assigned, conditional on covariates and unobserved effects. The additional restrictions in (2.3) are exogeneity conditions on $(\mathbf{X}^{NT}, \mathbf{A}^N, \mathbf{B}^T)$ with respect to U_{it} , conditional on \mathbf{A}_i and \mathbf{B}_t . The most substantive is the exogeneity of X_{it} . Given (2.2), this is a mild condition as time homogeneity already imposes that any relationship between U_{it} and X_{it} can only be unit and time-invariant. Taken together, (2.2) and (2.3) impose that

$$U_{it} \mid \mathbf{A}_i, \mathbf{B}_t \stackrel{d}{=} U_{js} \mid \mathbf{A}_j, \mathbf{B}_s, \quad \text{for all } i, j \in \mathbb{N}, t, s \in \mathbb{T}. \tag{2.4}$$

The product support condition guarantees overlap in the support of the unobserved effects for all values of the treatments. This condition is similar to the overlap condition used in cross section treatment effect models under unconfoundedness or selection on observables. Thus, together with (2.3), it implies that $P_{it}(x) := \Pr(X_{it} = x \mid \mathbf{A}^N, \mathbf{B}^T) > 0$, a.s., for all $i \in \mathbb{N}$, $t \in \mathbb{T}$ and $x \in \mathbb{X}$, where $P_{it}(x)$ is the analog of the propensity score in our setting. This condition is plausible in many applications; for example, in our empirical application in Section 5.1, $X_{it} = \mathbb{1}\{t \geq \tau_i\}$, where τ_i is the date of the law change in state i . In that case, if we consider τ_i to be a random variable with sufficiently large support conditional on the unobserved effects, then the condition $P_{it}(x) > 0$, a.s., is satisfied.

The model considered is similar to the static model in Chernozhukov et al. (2013), but there are three important differences. First, the structural function g has time effects as arguments and therefore allows the relationship between Y_{it} and X_{it} to vary over time in an unrestricted fashion even under (2.2). For example, it can include location and scale time effects. Second, Chernozhukov et al. (2013) impose that Y_{it} and X_{it} are identically distributed across i , which is stronger than the unit homogeneity in (2.4). Thus, unit homogeneity does not restrict the treatment assignment process. Third, they analyse short panels, whereas we rely on large T for identification. Our model also encompasses the nonseparable model with time effects in Freyberger (2017), where in our notation $Y_{it} = g_t(X_{it}, \mathbf{A}_i^T \mathbf{B}_t + U_{it})$.² We provide more examples of models covered by Assumption 2.1 below.

The structural function g is generally not identified, but can be used to construct interesting effects. Let $Y_{it}(\mathbf{x}) := g(\mathbf{x}, \mathbf{A}_i, \mathbf{B}_t, U_{it}(\mathbf{x}))$ be the potential outcome for individual i at time t obtained by setting exogenously $X_{it} = \mathbf{x} \in \mathbb{X}$, where

$$U_{it}(\mathbf{x}) \stackrel{d}{=} U_{it} \mid \mathbf{A}^N, \mathbf{B}^T. \tag{2.5}$$

¹ In the Aldous-Hoover representation, \mathbf{A}_i , \mathbf{B}_t , and U_{it} are independent uniform random variables.

² Note that our model allows for g to depend on t because the dimension of \mathbf{B}_t is unspecified.

Here we impose rank similarity as the distribution of $U_{it}(\mathbf{x})$ conditional on \mathbf{A}^N and \mathbf{B}^T does not change with \mathbf{x} . The main effects of interest are the average structural functions (ASFs)

$$\mu_t(\mathbf{x}) := \frac{1}{N} \sum_{i=1}^N E[Y_{it}(\mathbf{x}) \mid \mathbf{A}^N, \mathbf{B}^T], \quad \mu(\mathbf{x}) := \frac{1}{T} \sum_{t=1}^T \mu_t(\mathbf{x}), \tag{2.6}$$

and the conditional average structural functions (CASFs)

$$\begin{aligned} \mu_t(\mathbf{x} \mid \mathbb{X}_0) &:= \frac{1}{N_t(\mathbb{X}_0)} \sum_{i=1}^N \mathbb{1}\{X_{it} \in \mathbb{X}_0\} E[Y_{it}(\mathbf{x}) \mid \mathbf{A}^N, \mathbf{B}^T], \quad N_t(\mathbb{X}_0) = \sum_{i=1}^N \mathbb{1}\{X_{it} \in \mathbb{X}_0\}, \\ \mu(\mathbf{x} \mid \mathbb{X}_0) &:= \frac{1}{n(\mathbb{X}_0)} \sum_{t=1}^T N_t(\mathbb{X}_0) \mu_t(\mathbf{x} \mid \mathbb{X}_0), \quad n(\mathbb{X}_0) = \sum_{t=1}^T N_t(\mathbb{X}_0), \end{aligned} \tag{2.7}$$

where $\mathbb{X}_0 \subseteq \mathbb{X}$, provided that $n(\mathbb{X}_0) > 0$. The ASFs and CASFs correspond to averages of the potential outcome $Y_{it}(\mathbf{x})$ at a given time period or aggregated over the observed time periods. In both cases the average is over the cross sectional units in the observed sample or finite population. Infinite-population versions of the effects can be obtained by taking probability limits as $N \rightarrow \infty$. If X_{it} includes only a binary treatment, the ASFs and CASFs can be used to form treatment effects. For example, $\mu(1) - \mu(0)$ is the time-aggregated average treatment effect and $\mu_t(1 \mid \{1\}) - \mu_t(0 \mid \{1\})$ is the average treatment effect on the treated at time t . Distribution structural functions (DSFs) can be constructed analogously replacing $Y_{it}(\mathbf{x})$ by $\mathbb{1}\{Y_{it}(\mathbf{x}) \leq y\}$ in (2.6) and (2.7) for $y \in \mathbb{Y}$. Quantile effects can then be formed by taking left-inverses of the DSFs and taking differences. For example, the τ -quantile treatment effect at time t is $q_{t,\tau}(1) - q_{t,\tau}(0)$, where

$$q_{t,\tau}(\mathbf{x}) = \inf \left\{ y \in \mathbb{Y} : \frac{1}{N} \sum_{i=1}^N E[\mathbb{1}\{Y_{it}(\mathbf{x}) \leq y\} \mid \mathbf{A}^N, \mathbf{B}^T] \geq \tau \right\}.$$

We provide some examples of data generating processes that satisfy Assumption 2.1. The purpose is to show that Assumption 2.1 covers a great variety of models commonly used in empirical analysis. Our estimation methods are generic in that we do not need to specify the data generating process, beside satisfying Assumption 2.1. Of course, using more information about the data generating process would lead to more efficient estimators, but at the cost of robustness to model misspecification.

EXAMPLE 2.1. (LINEAR FACTOR MODEL) Consider the linear panel model with factor structure in the error terms:

$$Y_{it}(\mathbf{x}) = \mathbf{x}^T \boldsymbol{\beta} + \boldsymbol{\lambda}_i^T \mathbf{f}_t + \sigma_i(\mathbf{x}) \sigma_t(\mathbf{x}) U_{it}(\mathbf{x}), \quad U_{it}(\mathbf{x}) \mid \mathbf{X}^{NT}, \mathbf{A}^N, \mathbf{B}^T \sim i.i.d. F_U,$$

where $U_{it}(\mathbf{x})$ is a zero mean random variable with marginal distribution F_U , which does not depend on \mathbf{x} . This is special case of Assumption 2.1 with $Y_{it} = Y_{it}(\mathbf{X}_{it})$, $\mathbf{A}_i = (\boldsymbol{\lambda}_i, \{\sigma_i(\mathbf{x}) : \mathbf{x} \in \mathbb{X}\})$, $\mathbf{B}_t = (\mathbf{f}_t, \{\sigma_t(\mathbf{x}) : \mathbf{x} \in \mathbb{X}\})$, and $U_{it} = U_{it}(\mathbf{X}_{it})$. The average effect of changing the covariate from \mathbf{x}_0 to \mathbf{x}_1 at t is

$$\mu_t(\mathbf{x}_1) - \mu_t(\mathbf{x}_0) = \mu_t(\mathbf{x}_1 \mid \{\mathbf{x}_1\}) - \mu_t(\mathbf{x}_0 \mid \{\mathbf{x}_1\}) = (\mathbf{x}_1 - \mathbf{x}_0)^T \boldsymbol{\beta}.$$

A version of this model was considered by Kim and Oka (2014) to analyse the effect of unilateral divorce laws on divorce rates in the US. This model encompasses the standard

difference-in-differences model, $Y_{it}(\mathbf{x}) = \mathbf{x}^T \boldsymbol{\beta} + \lambda_i + f_t + \sigma_i(\mathbf{x})\sigma_t(\mathbf{x})U_{it}(\mathbf{x})$, by setting $\lambda_i = (\lambda_i, 1)^T$ and $\mathbf{f}_t = (1, f_t)^T$.

EXAMPLE 2.2. (BINARY RESPONSE MODEL) Assume that the potential outcome $Y_{it}(\mathbf{x})$ is binary and generated by

$$Y_{it}(\mathbf{x}) = \mathbb{1}\{m(\mathbf{x}, \mathbf{A}_i, \mathbf{B}_t) \geq U_{it}(\mathbf{x})\}, \quad U_{it}(\mathbf{x}) \mid \mathbf{X}^{NT}, \mathbf{A}^N, \mathbf{B}^T \sim i.i.d. \mathcal{U}(0, 1),$$

for some unknown function m . Here, assuming that $U_{it}(\mathbf{x})$ is uniform is a normalization, since m can be arbitrary. This latent index model with unobserved effects is a special case of Assumption 2.1 with $Y_{it} = Y_{it}(\mathbf{X}_{it})$ and $\mathbf{U}_{it} = U_{it}(\mathbf{X}_{it})$. The ASFs at \mathbf{x} and t is

$$\mu_t(\mathbf{x}) = \frac{1}{N} \sum_{i=1}^N m(\mathbf{x}, \mathbf{A}_i, \mathbf{B}_t).$$

Similar latent index models for count or censored responses are also covered by Assumption 2.1.

EXAMPLE 2.3. (TREATMENT EFFECT FACTOR MODEL) Assume that \mathbf{X}_{it} contains only a binary treatment indicator, i.e., $\mathbb{X} = \{0, 1\}$. The potential outcomes are generated by the linear factor model

$$Y_{it}(\mathbf{x}) = \lambda_i(\mathbf{x})^T \mathbf{f}_t(\mathbf{x}) + \sigma_i(\mathbf{x})\sigma_t(\mathbf{x})U_{it}(\mathbf{x}), \quad U_{it}(\mathbf{x}) \mid \mathbf{X}^{NT}, \mathbf{A}^N, \mathbf{B}^T \sim i.i.d. F_U, \quad \mathbf{x} \in \mathbb{X},$$

where $U_{it}(\mathbf{x})$ is a zero mean random variable with marginal distribution F_U , which does not depend on \mathbf{x} . This is special case of Assumption 2.1 with $Y_{it} = Y_{it}(\mathbf{X}_{it})$, $\mathbf{A}_i = (\{\lambda_i(\mathbf{x}), \sigma_i(\mathbf{x}) : \mathbf{x} \in \mathbb{X}\})$, $\mathbf{B}_t = (\{\mathbf{f}_t(\mathbf{x}), \sigma_t(\mathbf{x}) : \mathbf{x} \in \mathbb{X}\})$, and $\mathbf{U}_{it} = U_{it}(\mathbf{X}_{it})$. The average treatment effect at t is

$$\mu_t(1) - \mu_t(0) = \frac{1}{N} \sum_{i=1}^N [\lambda_i(1)^T \mathbf{f}_t(1) - \lambda_i(0)^T \mathbf{f}_t(0)],$$

and the average effect on the treated at t is

$$\mu_t(1 \mid \{1\}) - \mu_t(0 \mid \{1\}) = \frac{1}{N_t(1)} \sum_{i=1}^N \mathbb{1}\{\mathbf{X}_{it} = 1\} [\lambda_i(1)^T \mathbf{f}_t(1) - \lambda_i(0)^T \mathbf{f}_t(0)],$$

provided that $N_t(1) = \sum_{i=1}^N \mathbb{1}\{\mathbf{X}_{it} = 1\} > 0$. Versions of this model have been considered by Hsiao et al. (2012), Gobillon and Magnac (2016), Athey et al. (2021), Li and Bell (2017), Xu (2017), Li (2018), Bai and Ng (2019a), Xiong and Pelger (2019), and Chan and Kwok (2021). Example 2.1 is a special case with $\lambda_i(\mathbf{x})^T \mathbf{f}_t(\mathbf{x}) = \mathbf{x}^T \boldsymbol{\beta} + \lambda_i^T \mathbf{f}_t$.

Throughout this paper we use standard panel data notation, with the two panel dimensions being denoted by units i and time t . However, one could also consider pseudo-panel or network applications of our results, where the two panel dimensions are denoted by i and j , and Y_{ij} could, for example, be wage of worker i in firm j , consumption of member i in household j , a friendship indicator between individuals i and j , or the volume of trade from country i to country j . The existing literature on two-way heterogeneity in network models usually either makes stronger parametric assumptions than we impose here (e.g., Graham, 2017; Dzemski, 2019; Zeleneev, 2020; Chen et al., 2021) or uses stochastic blockmodels or graphon models, which typically ignore the effect of covariates (e.g., Holland et al., 1983; Wolfe and Olhede, 2013; Gao et al., 2015; Auerbach, 2019). Our methods of estimating nonparametric models with two-way heterogeneity may therefore also be of interest in a network context.

3. ESTIMATION VIA FACTOR STRUCTURE APPROXIMATION

A natural starting point to estimate the effects in (2.6) and (2.7) is to use empirical analogs. This amounts to replacing $E[Y_{it}(\mathbf{x}) | A^N, \mathbf{B}^T]$ by an estimator. There are two complications with this approach. First, the potential outcome $Y_{it}(\mathbf{x})$ is not observable. We deal with this complication by noting that

$$\begin{aligned} E[Y_{it}(\mathbf{x}) | A^N, \mathbf{B}^T] &= E[g(\mathbf{x}, \mathbf{A}_i, \mathbf{B}_t, U_{it}(\mathbf{x})) | A^N, \mathbf{B}^T] \\ &= E[g(\mathbf{x}, \mathbf{A}_i, \mathbf{B}_t, U_{it}) | A^N, \mathbf{B}^T] = E[g(\mathbf{x}, \mathbf{A}_i, \mathbf{B}_t, U_{it}) | X_{it} = \mathbf{x}, \mathbf{A}_i, \mathbf{B}_t] \\ &= E[Y_{it} | X_{it} = \mathbf{x}, \mathbf{A}_i, \mathbf{B}_t], \end{aligned}$$

under the rank similarity in (2.5) and Assumption 2.1. Hence, we can write the expectation of the potential outcome as an expectation of the observed outcome. The second complication is that \mathbf{A}_i and \mathbf{B}_t are not observable, so that we cannot directly estimate $E[Y_{it} | X_{it} = \mathbf{x}, \mathbf{A}_i, \mathbf{B}_t]$. To deal with this complication, we start by noticing that

$$\begin{aligned} E[Y_{it} | X_{it} = \mathbf{x}, \mathbf{A}_i = \mathbf{a}, \mathbf{B}_t = \mathbf{b}] &= E[g(\mathbf{x}, \mathbf{a}, \mathbf{b}, U_{it}) | \mathbf{A}_i = \mathbf{a}, \mathbf{B}_t = \mathbf{b}] \\ &=: m(\mathbf{x}, \mathbf{a}, \mathbf{b}), \end{aligned} \tag{3.1}$$

where the function m does not vary with i and t , by implication (2.4) of Assumption 2.1. We show next how this function can be approximated and estimated using a low-rank factor structure.

3.1. Low-rank factor structure approximation

For ease of exposition, we assume in the rest of the paper that the covariate vector X_{it} includes only a binary treatment and $\mathbb{X} = \{0, 1\}$. Accordingly, we denote the covariate and its values by X_{it} and x instead of \mathbf{X}_{it} and \mathbf{x} . In what follows, x denotes a generic element of \mathbb{X} and all the assumptions and results hold for all $x \in \mathbb{X}_1 \subseteq \mathbb{X}$, where $\mathbb{X}_1 = \mathbb{X}$ if we are interested in the entire population, $\mathbb{X}_1 = \{0\}$ if we are interested in the treated subpopulation, and $\mathbb{X}_1 = \{1\}$ if we are interested in the untreated subpopulation.

The approximation that we propose is based on the singular value decomposition of the function $(\mathbf{a}, \mathbf{b}) \mapsto m(x, \mathbf{a}, \mathbf{b})$ for each $x \in \mathbb{X}$. We make two assumptions on this decomposition. The first assumption is a sampling condition on the unobserved effects that will be useful to define a norm for the eigenfunctions.

ASSUMPTION 3.1. (SAMPLING OF \mathbf{A}_i and \mathbf{B}_t) (i) \mathbf{A}_i is independent and identically distributed across $i \in \mathbb{N}$ with distribution F_A , (ii) \mathbf{B}_t is independent and identically distributed over $t \in \mathbb{T}$ with distribution F_B , and (iii) \mathbf{A}_i and \mathbf{B}_t are independent for all i, t .

For simplicity we consider the case where both \mathbf{A}_i and \mathbf{B}_t are independently distributed across i and over t , but since we consider asymptotic sequences where both N and T become large one could also allow for appropriate weak dependence across both i and t . Formalizing this weak dependence would complicate both the assumption and the proof of the following results, which is why we decided to stick to independence in our presentation here.

The next assumption is a regularity condition on the function $m(x, \mathbf{a}, \mathbf{b})$.

ASSUMPTION 3.2. (SMOOTHNESS OF $(\mathbf{a}, \mathbf{b}) \mapsto m(x, \mathbf{a}, \mathbf{b})$) The function $(\mathbf{a}, \mathbf{b}) \mapsto m(x, \mathbf{a}, \mathbf{b})$ admits a singular value decomposition

$$m(x, \mathbf{a}, \mathbf{b}) = \sum_{j=1}^{\infty} s_j(x) u_j(x, \mathbf{a}) v_j(x, \mathbf{b}),$$

under the $L_2(F_A \times F_B)$ norm, where the eigenfunctions $u_j(x, \mathbf{a})$ and $v_j(x, \mathbf{b})$ are orthonormal, i.e.,

$$\begin{aligned} \mathbb{E} u_j(x, \mathbf{A}_i)^2 &= 1, & \mathbb{E} u_j(x, \mathbf{A}_i) u_k(x, \mathbf{A}_i) &= 0, \\ \mathbb{E} v_j(x, \mathbf{B}_t)^2 &= 1, & \mathbb{E} v_j(x, \mathbf{B}_t) v_k(x, \mathbf{B}_t) &= 0, & j \neq k \in \{1, 2, 3 \dots\}, \end{aligned}$$

and the singular values $s_1(x) \geq s_2(x) \geq s_3(x) \geq \dots \geq 0$ satisfy

$$\sum_{j=1}^{\infty} s_j(x) < \infty.$$

There is a large literature on singular value decompositions of functions, which shows that, under appropriate conditions, the singular values satisfy $s_j(x) \lesssim j^{-\alpha}$,³ where the decay coefficient α depends on the dimensions of the arguments \mathbf{a}, \mathbf{b} , and on the smoothness of $(\mathbf{a}, \mathbf{b}) \mapsto m(x, \mathbf{a}, \mathbf{b})$. For sufficiently smooth functions, $\alpha > 1$ and therefore $\sum_{j=1}^{\infty} s_j(x) < \infty$. For example, if $(\mathbf{a}, \mathbf{b}) \mapsto m(x, \mathbf{a}, \mathbf{b})$ is continuously differentiable up to order s and \mathbb{A} and \mathbb{B} are compact, then

$$s_j(x) \lesssim j^{-\frac{s}{d_a \wedge d_b}},$$

by theorem 3.3 of Griebel and Harbrecht (2013), where $d_a \wedge d_b$ is the minimum of d_a and d_b . This implies that $\sum_{j=1}^{\infty} s_j(x) < \infty$ if $s > d_a \wedge d_b$. Assumption 3.2 is therefore a high-level smoothness assumption on $(\mathbf{a}, \mathbf{b}) \mapsto m(x, \mathbf{a}, \mathbf{b})$, very similar to the assumption 2.2. in Menzel (2018), where an analogous condition on the singular values is imposed, with the same aim of controlling the behaviour of a function of unobserved two-dimensional heterogeneity.

The formulation of this smoothness assumption is convenient for our purposes, because it immediately leads to a low-rank approximation of $m(x, \mathbf{a}, \mathbf{b})$. The low-rank approximation truncates the singular value decomposition to the first R elements,

$$\begin{aligned} m(x, \mathbf{a}, \mathbf{b}) &= \sum_{j=1}^{\infty} \underbrace{s_j(x)^{1/2} u_j(x, \mathbf{a})}_{=: \phi_j(x, \mathbf{a})} \underbrace{s_j(x)^{1/2} v_j(x, \mathbf{b})}_{=: \psi_j(x, \mathbf{b})} \\ &= \sum_{j=1}^R \phi_j(x, \mathbf{a}) \psi_j(x, \mathbf{b}) + \zeta_R(x, \mathbf{a}, \mathbf{b}). \end{aligned} \tag{3.2}$$

The first term is the approximation and the second term is the approximation error. Under Assumption 3.2,

$$\mathbb{E} \zeta_R(x, \mathbf{A}_i, \mathbf{B}_t)^2 \rightarrow 0 \text{ as } R \rightarrow \infty.$$

³ Here, $s_j(x) \lesssim j^{-\alpha}$ means that there exists a constant $c > 0$ such that $s_j(x) \leq c j^{-\alpha}$, for all j .

In other words, the approximation error can be made negligible by increasing the truncation point R . For example, if $s_j(x) \lesssim j^{-\alpha}$ with $\alpha > 1$, then

$$\begin{aligned} E \zeta_R(x, \mathbf{A}_i, \mathbf{B}_t)^2 &= E \left[\sum_{j=R+1}^{\infty} s_j(x) u_j(x, \mathbf{A}_i) v_j(x, \mathbf{B}_t) \right]^2 \\ &= \sum_{j,k=R+1}^{\infty} s_j(x) s_k(x) E[u_j(x, \mathbf{A}_i) u_k(x, \mathbf{A}_i)] E[v_j(x, \mathbf{B}_t) v_k(x, \mathbf{B}_t)] \\ &= \sum_{j=R+1}^{\infty} s_j(x)^2 \lesssim \sum_{j=R+1}^{\infty} j^{-2\alpha} \leq \int_R^{\infty} j^{-2\alpha} dj \lesssim R^{1-2\alpha}, \end{aligned}$$

by Assumptions 3.1 and 3.2. Hence, $\zeta_R(x, \mathbf{A}_i, \mathbf{B}_t)$ converges in mean square to zero at a polynomial rate with R .

Combining (3.1) and (3.2), we obtain the approximate factor model

$$Y_{it} = \lambda_i(X_{it})^T \mathbf{f}_t(X_{it}) + \zeta_R(X_{it}, \mathbf{A}_i, \mathbf{B}_t) + E_{it}, \quad E_{it} := Y_{it} - E[Y_{it} | X_{it}, \mathbf{A}_i, \mathbf{B}_t], \quad (3.3)$$

where $\lambda_i(x) = [\phi_1(x, \mathbf{A}_i), \dots, \phi_R(x, \mathbf{A}_i)]^T$, $\mathbf{f}_t(x) = [\psi_1(x, \mathbf{B}_t), \dots, \psi_R(x, \mathbf{B}_t)]^T$, and the composite error $v_{it} := \zeta_R(X_{it}, \mathbf{A}_i, \mathbf{B}_t) + E_{it}$ contains the approximation error, $\zeta_R(X_{it}, \mathbf{A}_i, \mathbf{B}_t)$, and the conditional expectation error, E_{it} . The factor structure can be seen as a series or sieve approximation to the function $(\mathbf{a}, \mathbf{b}) \mapsto m(x, \mathbf{a}, \mathbf{b})$ with basis functions $\{\phi_j(x, \mathbf{a}) \psi_j(x, \mathbf{b})\}_{j=1}^{\infty}$ if we let $R = R_{N,T}$ to grow with N and T such that $\zeta_R(x, \mathbf{a}, \mathbf{b})$ vanishes as $N, T \rightarrow \infty$. The factor structure approximation is exact in some cases for fixed R . For instance, in Example 2.3

$$m(x, \mathbf{A}_i, \mathbf{B}_t) = \lambda_i(x)^T \mathbf{f}_t(x),$$

so that $\zeta_R(x, \mathbf{A}_i, \mathbf{B}_t) = 0$, a.s., if R is greater or equal to the number of factors.

In the model (3.3) the factor structure changes with the treatment level. In other words, we have a different pure factor model for each $x \in \mathbb{X}$, that is

$$Y_{it} = \lambda_i(x)^T \mathbf{f}_t(x) + v_{it} \text{ if } X_{it} = x.$$

This observation leads to our first estimation strategy where the data is partitioned by the treatment level and separate factors and factor loadings are estimated in each element of the partition by solving the least squares program

$$\min_{\{\lambda_i\}_{i=1}^N, \{\mathbf{f}_t\}_{t=1}^T} \frac{1}{2} \sum_{i=1}^N \sum_{t=1}^T D_{it}(x) (Y_{it} - \lambda_i^T \mathbf{f}_t)^2, \quad (3.4)$$

where $D_{it}(x) := \mathbb{1}\{X_{it} = x\}$. Unfortunately, we cannot solve this problem using standard principal component analysis due to the presence of missing data, that is, each observational unit (i, t) is not available at all treatment levels. In the next section, we apply matrix-completion methods to deal with this problem.

3.2. Estimation by matrix-completion methods

We start by expressing the program (3.4) in matrix form. Let $\mathbf{\Gamma}^R(x) = \lambda^N(x) \mathbf{f}^T(x)^T$, where $\lambda^N(x) = [\lambda_1(x), \dots, \lambda_N(x)]^T$, a $N \times R$ matrix of factor loadings, and $\mathbf{f}^T(x) =$

$[\mathbf{f}_1(x), \dots, \mathbf{f}_T(x)]^T$, a $T \times R$ matrix of factors. The least squares estimator of $\mathbf{\Gamma}^R(x)$ is the $N \times T$ matrix $\mathbf{\Gamma}$ with typical element Γ_{it} that solves

$$\min_{\{\mathbf{\Gamma} \in \mathbb{R}^{N \times T} : \text{rank}(\mathbf{\Gamma}) \leq R\}} \frac{1}{2} \sum_{i=1}^N \sum_{t=1}^T D_{it}(x) (Y_{it} - \Gamma_{it})^2. \tag{3.5}$$

Let $\mathbf{Y}(x)$ be a $N \times T$ matrix whose (i, t) element is Y_{it} if $X_{it} = x$ and is missing otherwise. The previous program is closely related to the problem of completing the missing entries of $\mathbf{Y}(x)$ using a low-rank approximation matrix $\mathbf{\Gamma}^R(x)$ (Rennie and Srebro, 2005; Candès and Recht, 2009; Candès and Tao, 2010). This connection was previously noticed by Athey et al. (2021) and Amjad et al. (2018) in the context of treatment effects models. The solution is the $N \times T$ matrix of rank R whose entries are the closest in the mean squared error sense to the corresponding entries of $\mathbf{Y}(x)$.

The previous program is combinatorially hard because of the constraint in the rank of the matrix Srebro and Jaakkola (2003). Following Fazel (2003) we consider the convex relaxation of this program. Let $\|\mathbf{M}\|_\infty$ be the spectral norm of a $\mathbb{R}^{N \times T}$ -matrix \mathbf{M} , and define the nuclear norm (also called trace norm) of $\mathbf{\Gamma}$ as the corresponding dual norm $\|\mathbf{\Gamma}\|_1 := \max_{\{\mathbf{M} \in \mathbb{R}^{N \times T} : \|\mathbf{M}\|_\infty \leq 1\}} \text{Tr}(\mathbf{M}\mathbf{\Gamma})$. This nuclear norm can equivalently be defined as the sum of the singular values of $\mathbf{\Gamma}$. Using this norm we can write the convex relaxation of the program (3.5) as follows,

$$\min_{\{\mathbf{\Gamma} \in \mathbb{R}^{N \times T} : \|\mathbf{\Gamma}\|_1 \leq R_1\}} \frac{1}{2} \sum_{i=1}^N \sum_{t=1}^T D_{it}(x) (Y_{it} - \Gamma_{it})^2,$$

where R_1 is a positive constant such that $R = f(R_1)$, where f is an increasing function. Hence, $\zeta_R(x, \mathbf{A}_i, \mathbf{B}_t)$ vanishes in mean square as $R_1 \rightarrow \infty$. We replace the rank constraint, $\text{rank}(\mathbf{\Gamma}) \leq R$, by a constraint on the nuclear norm of the matrix, $\|\mathbf{\Gamma}\|_1 \leq R_1$, i.e., we replace a constraint in the number of nonzero singular values by a constraint in the sum of singular values. This program is convex in $\mathbf{\Gamma}$ and can be reformulated in Lagrange form as

$$\min_{\{\mathbf{\Gamma} \in \mathbb{R}^{N \times T}\}} \frac{1}{2} \sum_{i=1}^N \sum_{t=1}^T D_{it}(x) (Y_{it} - \Gamma_{it})^2 + \rho(R_1) \|\mathbf{\Gamma}\|_1, \tag{3.6}$$

where $\rho(R_1) \geq 0$ is a regularization parameter, which is a one-to-one increasing function of R_1 . There exist efficient algorithms to solve this program (Mazumder et al., 2010).

Let $\hat{\mathbf{\Gamma}}(x)$ be a solution to (3.6) with typical element $\hat{\Gamma}_{it}(x)$. Then, we can form estimators of the ASF and CASF as

$$\hat{\mu}_i(x) = \frac{1}{N} \sum_{i=1}^N [D_{it}(x)Y_{it} + \{1 - D_{it}(x)\}\hat{\Gamma}_{it}(x)],$$

and

$$\hat{\mu}_i(x | \{x_0\}) = \frac{\sum_{i=1}^N D_{it}(x_0) [D_{it}(x)Y_{it} + \{1 - D_{it}(x)\}\hat{\Gamma}_{it}(x)]}{\sum_{i=1}^N D_{it}(x_0)}.$$

In the next section, we provide conditions under which these estimators are consistent using asymptotic sequences where $N, T \rightarrow \infty$. These estimators, however, might display shrinkage biases in finite samples due to the nuclear norm regularization (Cai et al., 2010; Ma et al., 2011; Bai and Ng, 2019b). We propose two matching procedures to debias the estimator in Section 4.

3.3. Consistency of matrix-completion estimator

Let $\Gamma^\infty(x)$ be the $N \times T$ matrix with typical element $\Gamma_{it}^\infty(x) = m(x, \mathbf{A}_i, \mathbf{B}_t)$ and $\mathbf{E}(x)$ be the $N \times T$ matrix with typical element

$$E_{it}(x) := \begin{cases} Y_{it} - \Gamma_{it}^\infty(x) & \text{if } X_{it} = x, \\ 0 & \text{otherwise.} \end{cases} \tag{3.7}$$

Note that $\Gamma^\infty(x) = \lim_{R \rightarrow \infty} \Gamma^R(x)$ a.s. Furthermore, we introduce the notation $\mathbb{D}(x) = \{(i, t) \in \mathbb{N} \times \mathbb{T} : X_{it} = x\}$, and $n(x) = |\mathbb{D}(x)|$ for the number of observations with $X_{it} = x$.

Recall that

$$\widehat{\Gamma}(x) \in \operatorname{argmin}_{\Gamma \in \mathbb{R}^{N \times T}} Q_{NT}(\Gamma, \rho, x), \quad Q_{NT}(\Gamma, \rho, x) = \frac{1}{2} \sum_{(i,t) \in \mathbb{D}(x)} (Y_{it} - \Gamma_{it})^2 + \rho \|\Gamma\|_1, \tag{3.8}$$

where $\rho := \rho(R_1)$. Here, if the argmin over $\Gamma \in \mathbb{R}^{N \times T}$ is not unique, then we can choose $\widehat{\Gamma}(x)$ arbitrarily from the set of minimizers—our results are not affected by that, we only require that $Q_{NT}(\widehat{\Gamma}(x), \rho, x) \leq Q_{NT}(\Gamma, \rho, x)$, for all $\Gamma \in \mathbb{R}^{N \times T}$. We want to show that $\widehat{\Gamma}(x)$ converges to $\Gamma^\infty(x)$ as $N, T \rightarrow \infty$ in some sense such that $\widehat{\mu}(x) - \mu(x) = o_P(1)$. For that we require additional assumptions.

ASSUMPTION 3.3. (ERROR MOMENTS) *Conditional on $\mathbf{X}^{NT}, \mathbf{A}^N$, and \mathbf{B}^T , $E_{it}(x)$ is independent across $(i, t) \in \mathbb{D}(x)$, and there exists a constant $b < \infty$ that does not depend on i, t, N, T , such that*

$$\mathbb{E} [E_{it}(x)^4 \mid \mathbf{A}^N, \mathbf{B}^T, \mathbf{X}^{NT}] \leq b.$$

Furthermore, we assume that $n(x)^{-1} \sum_{(i,t) \in \mathbb{D}(x)} \Gamma_{it}^\infty(x)^2 = O_P(1)$.

For the purpose of showing Lemma 3.1 and Theorem 3.1, we could alternatively replace Assumption 3.3 by the two high-level conditions:

$$\frac{2}{n(x)} \sum_{(i,t) \in \mathbb{D}(x)} \Gamma_{it}^\infty(x) E_{it} = o_P(1), \quad \|\mathbf{E}(x)\|_\infty = O_P(\sqrt{N+T}),$$

where again $\|\cdot\|_\infty$ denotes the spectral norm. The first of those conditions is implied by Assumption 3.3 through application of the weak law of large numbers, while the second follows, for example, by the spectral norm inequality in Latała (2005). In principle, we could still derive those high-level conditions if we allowed for appropriate weak dependence of $E_{it}(x)$ across i and over t , but we again focus on the independent case for simplicity of presentation.

We first provide a consistency result for the entries of $\widehat{\Gamma}(x)$ that correspond to the observed values of $\mathbf{Y}(x)$.

LEMMA 3.1. *Let the Assumptions 3.1, 3.2, and 3.3 hold, and assume that $\rho = \rho_{NT}$ is chosen such that $\rho_{NT}/\sqrt{N+T} \rightarrow \infty$ and $\rho_{NT}\sqrt{NT}/n(x) \rightarrow 0$ as $N, T \rightarrow \infty$. Then,*

$$\frac{1}{n(x)} \sum_{(i,t) \in \mathbb{D}(x)} [\widehat{\Gamma}_{it}(x) - \Gamma_{it}^\infty(x)]^2 = o_P(1).$$

A necessary condition for the existence of the sequence $\rho = \rho_{NT}$ in Lemma 3.1 is $n(x)/\sqrt{(N+T)NT} \rightarrow \infty$; that is, the fraction $n(x)/(NT)$ of observations with $X_{it} = x$ can converge to zero, but not too fast. Apart from that, Lemma 3.1 does not restrict the assignment

process that determines X^{NT} . Notice also that Lemma 3.1 does not require Assumption 2.1 because $\Gamma^\infty(x)$ is a reduced-form parameter.

Applying the Cauchy-Schwarz inequality

$$\left(\frac{1}{n(x)} \sum_{(i,t) \in \mathbb{D}(x)} a_{it} \right)^2 \leq \frac{1}{n(x)} \sum_{(i,t) \in \mathbb{D}(x)} a_{it}^2,$$

with $a_{it} = \widehat{\Gamma}_{it}(x) - \Gamma_{it}^\infty(x)$, Lemma 3.1 guarantees that

$$\frac{1}{n(x)} \sum_{(i,t) \in \mathbb{D}(x)} [\widehat{\Gamma}_{it}(x) - \Gamma_{it}^\infty(x)] = o_P(1).$$

Nevertheless, Lemma 3.1 is not directly useful to show the consistency of the estimators of the ASF, because it only guarantees L_2 -consistency of $\widehat{\Gamma}(x)$ over the set of entries (i, t) for which $X_{it} = x$. Those are exactly the observations for which an unbiased estimator of $\Gamma_{it}^\infty(x) = m(x, \mathbf{A}_i, \mathbf{B}_t)$ is already available, namely Y_{it} . The consistency result we would like to obtain is

$$\frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T [\widehat{\Gamma}_{it}(x) - \Gamma_{it}^\infty(x)]^2 = o_P(1), \tag{3.9}$$

but such a result will certainly require stronger assumptions on X^{NT} than we have imposed so far.

The existing literature on matrix completion relies on the concept of *restricted strong convexity* to derive (3.9). This approach shows that under certain conditions on a $\mathbb{R}^{N \times T}$ -matrix \mathbf{M} with entries M_{it} , and on X^{NT} (which determines the set $\mathbb{D}(x)$), there exists a constant $c > 0$ such that with high probability

$$\frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T M_{it}^2 \leq \frac{c}{n(x)} \sum_{(i,t) \in \mathbb{D}(x)} M_{it}^2.$$

See theorem 1 in Negahban and Wainwright (2012), lemma 12 in Klopp et al. (2014), and lemma 3 in Athey et al. (2021). Thus, if $M_{it} = \widehat{\Gamma}_{it}(x) - \Gamma_{it}^\infty(x)$ and X^{NT} satisfy restricted strong convexity, then (3.9) would follow from Lemma 3.1.

We pursue a different strategy than the existing matrix completion literature to show that

$$\widehat{\mu}(x) := \frac{1}{T} \sum_{t=1}^T \widehat{\mu}_t(x) = \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T D_{it}(x) Y_{it} + \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T [1 - D_{it}(x)] \widehat{\Gamma}_{it}(x),$$

is a consistent estimator of $(NT)^{-1} \sum_{i=1}^N \sum_{t=1}^T \Gamma_{it}^\infty$, which under Assumption 2.1 is equal to $\mu(x)$ defined in (2.6). We believe that our approach is simpler in the setting of this paper where $\Gamma_{it}^\infty(x)$ is not necessarily of low rank. In particular, we do not aim to show (3.9), but instead we derive consistency of $\widehat{\mu}(x)$ directly. However, the following theorem still requires additional assumptions on the assignment process that determines X^{NT} , in the same way that additional conditions on X^{NT} are required to verify restricted strong convexity. For simplicity, we focus on consistency of $\widehat{\mu}(x)$ in the main text, but results for more general weighted averages of the form $(NT)^{-1} \sum_{i=1}^N \sum_{t=1}^T W_{it}(x) \Gamma_{it}^\infty(x)$, with known weights $W_{it}(x) \in \mathbb{R}$, are presented in Appendix A. For example, in the case of the treatment effects on the treated that we consider in the empirical application of Section 5.1, $W_{it}(x) = n(1)^{-1} NT X_{it}$.

THEOREM 3.1. *Let the Assumptions 2.1, 3.1, 3.2, and 3.3 hold. Consider $N, T \rightarrow \infty$ at the same rate, and let $\rho = \rho_{NT}$ be chosen such that $\rho_{NT}/\sqrt{N+T} \rightarrow \infty$ and $\rho_{NT}/\sqrt{NT} \rightarrow 0$. Let $P_{it}(x) = \Pr(X_{it} = x \mid \mathbf{A}^N, \mathbf{B}^T)$, and assume that $(NT)^{-1} \sum_{i=1}^N \sum_{t=1}^T P_{it}^{-1}(x) = O_P(1)$. Let $\mathbf{G}(x)$ be the $N \times T$ matrix with entries $G_{it}(x) = P_{it}^{-1}(x)(D_{it}(x) - P_{it}(x))$, and assume that $\|\mathbf{G}(x)\|_\infty = O_P(\sqrt{N+T})$, and*

$$\frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T P_{it}^{-1}(x) G_{it}(x) = o_P(1), \quad \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \Gamma_{it}^\infty(x) G_{it}(x) = o_P(1). \tag{3.10}$$

Then,

$$\widehat{\mu}(x) = \mu(x) + o_P(1).$$

To interpret the conditions in Theorem 3.1, notice that due to the definitions $D_{it}(x) = \mathbb{1}\{X_{it} = x\}$ and $P_{it}(x) = \Pr(X_{it} = x \mid \mathbf{A}^N, \mathbf{B}^T)$, $E[G_{it}(x) \mid \mathbf{A}^N, \mathbf{B}^T] = 0$ by construction, and $G_{it}(x)$ therefore plays a role very similar to the error term $E_{it}(x)$. In particular, the conditions in (3.10) can be verified by a weak law of large numbers, as long as $P_{it}^{-1}(x)$ is not too large, and $G_{it}(x)$ is not too strongly correlated across i and over t . Regarding the condition on the spectral norm $\|\mathbf{G}(x)\|_\infty = O_P(\sqrt{N+T})$, there are many results in the random-matrix theory literature that show this rate for mean-zero random matrices $\mathbf{G}(x)$; see, for example, Geman (1980), Bai et al. (1988), Silverstein (1989), and Yin et al. (1988). In particular, if $G_{it}(x)$ is independent across both i and t , then this rate result follows from the very elegant spectral norm inequality in Latała (2005), see the proof of Lemma 3.1 in Appendix A, where we apply that inequality to $E_{it}(x)$. However, that simple argument would require X_{it} to be independently distributed across i and t , conditional on $\mathbf{A}^N, \mathbf{B}^T$. More generally, we expect $\|\mathbf{G}(x)\|_\infty = O_P(\sqrt{N+T})$ to hold whenever the matrix entries $G_{it}(x)$ have zero mean, sufficiently bounded moments, and weak correlation across both i and t , see section S.2 of the supplementary material of Moon and Weidner (2017) for details.

We have thus shown that consistent estimates for ASFs can be obtained via the matrix-completion estimator even if the estimand $\Gamma_{it}^\infty(x) = m(x, \mathbf{A}_i, \mathbf{B}_t)$ itself is not of low rank. This is the main technical result of this paper. However, inference on $\mu(x)$ based on $\widehat{\mu}(x)$ can be problematic, because $\widehat{\mu}(x)$ is subject to both low-rank approximation and shrinkage biases. The low-rank approximation bias is due to the approximation error $\zeta_R(x, \mathbf{a}, \mathbf{b})$ in the decomposition of $m(x, \mathbf{a}, \mathbf{b})$ in equation (3.2). The shrinkage bias comes from bias in $\widehat{\Gamma}(x)$ due to the presence of the nuclear norm penalization in the objective function of (3.8). To isolate this bias, consider a simple case where $Y_{it}(x)$ follows a deterministic pure factor model

$$Y_{it}(x) = \Gamma_{it}(x) = \sum_{j=1}^R s_j(x) u_j(x, \mathbf{A}_i) v_j(x, \mathbf{B}_t).$$

Then, the matrix-completion estimator of $\Gamma_{it}(x)$ in (3.8) yields

$$\widehat{\Gamma}_{it}(x) = \sum_{j=1}^R [s_j(x) - \rho]_+ u_j(x, \mathbf{A}_i) v_j(x, \mathbf{B}_t),$$

where $[z]_+ = \max(z, 0)$. Compared to $\Gamma(x)$, $\widehat{\Gamma}(x)$ has the same eigenvectors but the singular values are shrunk toward zero. This argument carries over to the case where $Y_{it}(x)$ follows an approximate factor structure (Cai et al., 2010; Ma et al., 2011; Bai and Ng, 2019b). Because of these biases, we explore alternative estimates for $\mu(x)$ in Section 4.

3.4. Covariates and fixed effects

As we mentioned in Section 2, exogenous covariates can be incorporated by conditioning on their values. This method can produce very noisy estimators in small samples unless the covariates take only on few values. Here we consider a semiparametric version of the model that imposes additivity in the effect of the exogenous covariates, which may be continuous, discrete, or mixed. It also allows for additive unobserved individual and time effects that might vary across the covariate level x . These effects can be subsumed in the factor structure, but are usually considered separately in empirical analysis as the estimators perform better without regularizing them (Athey et al., 2021).

Let C_{it} be a d_c -vector of covariates, $\alpha(x) = (\alpha_1(x), \dots, \alpha_N(x))$ be a N -vector of individual effects and $\delta(x) = (\delta_1(x), \dots, \delta_T(x))$ be a T -vector of time effects. Then, we can replace the program (3.6) by

$$\min_{\{\beta \in \mathbb{R}^{d_c}, \alpha \in \mathbb{R}^N, \delta \in \mathbb{R}^T, \Gamma \in \mathbb{R}^{N \times T}\}} \sum_{i=1}^N \sum_{t=1}^T \mathbb{1}\{X_{it} = x\} (Y_{it} - C_{it}^T \beta - \alpha_i - \delta_t - \Gamma_{it})^2 + \rho(R_1) \|\Gamma\|_1,$$

Chernozhukov et al. (2018), Moon and Weidner (2018), and Beyhum and Gautier (2019) provide algorithms to solve this program. Let $\hat{\beta}(x)$, $\hat{\alpha}(x) = (\hat{\alpha}_1(x), \dots, \hat{\alpha}_N(x))$, $\hat{\delta}(x) = (\hat{\delta}_1(x), \dots, \hat{\delta}_T(x))$, and $\hat{\Gamma}(x)$ be the solution of the previous program. We can form estimators of the ASF and CASF as

$$\hat{\mu}_t(x) = \frac{1}{N} \sum_{i=1}^N [\mathbb{1}\{X_{it} = x\} Y_{it} + \mathbb{1}\{X_{it} \neq x\} \{C_{it}^T \hat{\beta}(x) + \hat{\alpha}_i(x) + \hat{\delta}_t(x) + \hat{\Gamma}_{it}(x)\}],$$

and

$$\hat{\mu}_t(x | \{x_0\}) = \frac{\sum_{i=1}^N [\mathbb{1}\{X_{it} = x_0 = x\} Y_{it} + \mathbb{1}\{X_{it} = x_0 \neq x\} \{C_{it}^T \hat{\beta}(x) + \hat{\alpha}_i(x) + \hat{\delta}_t(x) + \hat{\Gamma}_{it}(x)\}]}{\sum_{i=1}^N \mathbb{1}\{X_{it} = x_0\}}.$$

4. DEBIASING USING MATCHING METHODS

The matrix-completion estimator of the ASF is generally biased. As we explained in Section 3.3, the bias comes from two sources: low-rank approximation bias and shrinkage bias. One could attempt to correct the shrinkage bias by shifting the singular values of $\hat{\Gamma}(x)$ upwards. However, inference results on the ASFs based on matrix completion are generally very difficult to obtain even if $\Gamma^\infty(x)$ is truly low rank. In our setting, the presence of the additional low-rank approximation bias makes this even more challenging. We instead discuss alternative estimators and show that they have significantly lower biases than the matrix-completion estimators in the numerical simulations of Section 5.2.

To construct the estimators of $\Gamma^\infty(x)$, we start by extracting the factor structure of $\hat{\Gamma}(x)$ in (3.8). Let $\hat{\lambda}_i(x)$ and $\hat{f}_i(x)$ be the $R \times 1$ vectors that satisfy

$$\hat{\Gamma}_{it}(x) = \hat{\lambda}_i(x)^T \hat{f}_t(x),$$

subject to the usual normalizations that $T^{-1} \sum_{t=1}^T \widehat{\mathbf{f}}_t(x) \widehat{\mathbf{f}}_t(x)^\top$ is the identity matrix of size R and $N^{-1} \sum_{i=1}^N \widehat{\boldsymbol{\lambda}}_i(x) \widehat{\boldsymbol{\lambda}}_i(x)^\top$ is a diagonal matrix. Next, we apply a matching procedure to this factor structure. In its simplest version, we estimate each entry $\Gamma_{it}^\infty(x)$ such that $X_{it} \neq x$, by matching with the observation with $X_{js} = x$ that is the nearest neighbour in terms of the vectors $\widehat{\boldsymbol{\lambda}}_i(x)$ and $\widehat{\mathbf{f}}_t(x)$. In particular, $\check{\Gamma}_{it}(x) = Y_{i^{**}(i,t,x), t^{**}(i,t,x)}$ where $i^{**}(i, t, x) \in \mathbb{N}$ and $t^{**}(i, t, x) \in \mathbb{T}$ are a solution to the program

$$\begin{aligned} \min_{j \in \mathbb{N}, s \in \mathbb{T}} & \|\widehat{\boldsymbol{\lambda}}_i(x) - \widehat{\boldsymbol{\lambda}}_j(x)\|^2 + \|\widehat{\mathbf{f}}_t(x) - \widehat{\mathbf{f}}_s(x)\|^2 \\ \text{s.t.} & \quad X_{js} = x. \end{aligned}$$

We also consider a two-way matching procedure that combines matching with a difference-in-differences approach. It consists of two steps:

- (i) For all $x \in \mathcal{X}$ and $(i, t) \in \mathbb{N} \times \mathbb{T}$ such that $X_{it} \neq x$, find the matches $i^*(i, t, x) \in \mathbb{N}$ and $t^*(i, t, x) \in \mathbb{T}$ that solve the program

$$\begin{aligned} \min_{j \in \mathbb{N}, s \in \mathbb{T}} & \|\widehat{\boldsymbol{\lambda}}_i(x) - \widehat{\boldsymbol{\lambda}}_j(x)\|^2 + \|\widehat{\mathbf{f}}_t(x) - \widehat{\mathbf{f}}_s(x)\|^2 \\ \text{s.t.} & \quad X_{is} = X_{jt} = X_{js} = x. \end{aligned}$$

- (ii) Estimate $\Gamma_{it}(x)$ by

$$\widetilde{\Gamma}_{it}(x) = Y_{i, t^*(i,t,x)} + Y_{i^*(i,t,x), t} - Y_{i^*(i,t,x), t^*(i,t,x)}.$$

In other words, we find the match (j, s) with $X_{js} = x$ that not only is the closest to (i, t) in terms of the estimated factor structure, but also corresponds to a unit j with $X_{jt} = x$ and a time period s with $X_{is} = x$. Then, we estimate the counterfactual $\Gamma_{it}(x)$ as a linear combination of Y_{jt} , Y_{is} , and Y_{js} .

The additional difference-in-differences step in the two-way procedure is useful to reduce bias. To see this, we can compare $\widetilde{\Gamma}_{it}(x)$ with the simple matching estimator $\check{\Gamma}_{it}(x)$. Thus, abstracting from the estimation error in the factors and loadings,

$$\begin{aligned} E[\check{\Gamma}_{it}(x) - \Gamma_{it}(x) \mid \mathbf{A}^N, \mathbf{B}^T, \mathbf{X}^{NT}] &= m(x, \mathbf{A}_{i^{**}(i,t,x)}, \mathbf{B}_{t^{**}(i,t,x)}) - m(x, \mathbf{A}_i, \mathbf{B}_t) \\ &= \mathcal{O}_P(\|\mathbf{A}_{i^{**}(i,t,x)} - \mathbf{A}_i\| + \|\mathbf{B}_{t^{**}(i,t,x)} - \mathbf{B}_t\|), \end{aligned}$$

by a first-order Taylor expansion of $(\mathbf{a}_i, \mathbf{b}_t) \mapsto m(x, \mathbf{a}_i, \mathbf{b}_t)$ around $(\mathbf{A}_i, \mathbf{B}_t)$; whereas

$$\begin{aligned} E[\widetilde{\Gamma}_{it}(x) - \Gamma_{it}(x) \mid \mathbf{A}^N, \mathbf{B}^T, \mathbf{X}^{NT}] &= m(x, \mathbf{A}_{i^*(i,t,x)}, \mathbf{B}_{t^*(i,t,x)}) - m(x, \mathbf{A}_i, \mathbf{B}_t) \\ &= \mathcal{O}_P(\|\mathbf{A}_{i^*(i,t,x)} - \mathbf{A}_i\|^2 + \|\mathbf{B}_{t^*(i,t,x)} - \mathbf{B}_t\|^2), \end{aligned}$$

by a second-order Taylor expansion of $(\mathbf{a}_i, \mathbf{b}_t) \mapsto m(x, \mathbf{a}_i, \mathbf{b}_t)$ around $(\mathbf{A}_i, \mathbf{B}_t)$. The two-way matching removes the leading term of the Taylor expansion, reducing the bias of the matching by one order of magnitude because $i^{**}(i, t, x) \neq i$ or $t^{**}(i, t, x) \neq t$. On the other hand, $\|\mathbf{A}_{i^*(i,t,x)} - \mathbf{A}_i\| \geq \|\mathbf{A}_{i^{**}(i,t,x)} - \mathbf{A}_i\|$ and $\|\mathbf{B}_{t^*(i,t,x)} - \mathbf{B}_t\| \geq \|\mathbf{B}_{t^{**}(i,t,x)} - \mathbf{B}_t\|$ a.s. because the two-way procedure imposes the additional restrictions $X_{is} = X_{jt} = x$. Whether the first or second order bias dominates would generally be determined by the proportion of observations with $X_{js} = x$ and the distributions of \mathbf{A}_i and \mathbf{B}_t . We provide a numerical comparison of the biases of the matching estimators in Section 5.2.

We develop the theory for a debiased estimator that allows for multiple matches and estimated factors and loadings. Multiple matches are expected to reduce dispersion at the cost of increasing bias. Let $\lambda_i = \lambda(x, A_i)$ and $f_t = f(x, B_t)$ be the transformations of A_i and B_t that are consistently estimated by $\hat{\lambda}_i$ and \hat{f}_t .⁴ We define

$$\mathbb{N}_i = \left\{ j \in \mathbb{N} \setminus \{i\} : \left\| \hat{\lambda}_i - \hat{\lambda}_j \right\| \leq \tau_{NT} \right\}, \quad \mathbb{T}_t = \left\{ s \in \mathbb{T} \setminus \{t\} : \left\| \hat{f}_t - \hat{f}_s \right\| \leq \nu_{NT} \right\},$$

for some bandwidth parameters $\tau_{NT} > 0$ and $\nu_{NT} > 0$. The debiased estimator of $\mu(x)$ is then given by

$$\tilde{\mu}(x) = \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \tilde{Y}_{it}(x),$$

with

$$\tilde{Y}_{it}(x) = \begin{cases} Y_{it} & \text{if } X_{it} = x, \\ \frac{1}{n_{it}} \sum_{j \in \mathbb{N}_i} \sum_{s \in \mathbb{T}_t} \mathbb{1}_{\{X_{is} = X_{jt} = X_{js} = x\}} (Y_{is} + Y_{jt} - Y_{js}) & \text{if } X_{it} \neq x \text{ and } n_{it} > 0, \\ \frac{1}{n(x)} \sum_{(j,s) \in \mathbb{D}(x)} Y_{js} & \text{if } n_{it} = 0, \end{cases} \quad (4.1)$$

where $n_{it} := \sum_{j \in \mathbb{N}_i} \sum_{s \in \mathbb{T}_t} \mathbb{1}_{\{X_{is} = X_{jt} = X_{js} = x\}}$. Here, for $X_{it} \neq x$, we construct the counterfactual $\tilde{Y}_{it}(x)$ by averaging over all units $(j, s) \in \mathbb{N}_i \times \mathbb{T}_t$ that satisfy the constraint $X_{is} = X_{jt} = X_{js} = x$. Notice that if $X_{it} \neq x$ and $n_{it} = 0$, then we cannot construct a suitable counterfactual by that method. In that case we assign $\tilde{Y}_{it}(x)$ the average of the observations with $X_{js} = x$ to make sure that $\tilde{\mu}(x)$ is always well-defined, but our assumption below guarantees that this rarely happens.

This estimator has similar debiasing properties to the nearest neighbour described above, but it is more tractable theoretically because it varies more smoothly with respect to the factors and loadings.

Indeed, $\tilde{\mu}(x)$ can be written as

$$\tilde{\mu}(x) = \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \omega_{it} Y_{it},$$

where the weights ω_{it} are functions of $\hat{\lambda}_j$ and \hat{f}_s for all $j \in \mathbb{N}$ and $s \in \mathbb{T}$. To show that $\tilde{\mu}(x)$ is a consistent estimator of $\mu(x)$, we use the following assumption:

ASSUMPTION 4.1. (TWO-WAY MATCHING ESTIMATOR) *There exists a sequence $\xi_{NT} > 0$ such that $\xi_{NT} \rightarrow 0$ as $N, T \rightarrow \infty$, and*

- (i) $\frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \mathbb{1}_{\{X_{it} \neq x \ \& \ n_{it} = 0\}} = O_P(\xi_{NT})$.
- (ii) Y_{it} is uniformly bounded over i, t, N, T .
- (iii) Y_{it} is independent across both i and t , conditional on X^{NT}, A^N, B^T .

⁴ The matching method discussed here is also applicable to settings where the matching is based on variables other than the estimated factor structure. These include, for example, cross section and time series averages of the observable variables. See Appendix A for a more general treatment.

- (iv) The function $(\mathbf{a}, \mathbf{b}) \mapsto m(x, \mathbf{a}, \mathbf{b})$ is at least twice continuously differentiable with uniformly bounded second derivatives.
- (v) There exists $c > 0$ such that $\|\mathbf{a}_1 - \mathbf{a}_2\| \leq c \|\boldsymbol{\lambda}(\mathbf{a}_1) - \boldsymbol{\lambda}(\mathbf{a}_2)\|$ for all $\mathbf{a}_1, \mathbf{a}_2 \in \mathbb{A}$, and $\|\mathbf{b}_1 - \mathbf{b}_2\| \leq c \|\mathbf{f}(\mathbf{b}_1) - \mathbf{f}(\mathbf{b}_2)\|$ for all $\mathbf{b}_1, \mathbf{b}_2 \in \mathbb{B}$.
- (vi) $\frac{1}{N} \sum_{i=1}^N \left(\|\widehat{\boldsymbol{\lambda}}_i - \boldsymbol{\lambda}_i\|^2 + \max_{j \in \mathbb{N}_i} \|\widehat{\boldsymbol{\lambda}}_j - \boldsymbol{\lambda}_j\|^2 \right) = O_P(\xi_{NT})$.
 $\frac{1}{T} \sum_{t=1}^T \left(\|\widehat{\mathbf{f}}_t - \mathbf{f}_t\|^2 + \max_{s \in \mathbb{T}_t} \|\widehat{\mathbf{f}}_s - \mathbf{f}_s\|^2 \right) = O_P(\xi_{NT})$.
- (vii) $\tau_{NT}^2 = O_P(\xi_{NT})$ and $\nu_{NT}^2 = O_P(\xi_{NT})$.
- (viii) $\frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T E[\omega_{it}^2 | \mathbf{X}^{NT}, \mathbf{A}^N, \mathbf{B}^T] = O_P(NT \xi_{NT}^2)$.
- (ix) Let $\mathbf{Y}_{-(i,t),-(j,s)}^{NT}$ be the outcome matrix \mathbf{Y}^{NT} , but with Y_{it} and Y_{js} replace by zero (or some other nonrandom number), and all other outcomes unchanged. We assume

$$\frac{1}{(NT)^2} \sum_{i,j=1}^N \sum_{t,s=1}^T \mathbb{1}\{(i,t) \neq (j,s)\} E \left[\left| \omega_{it}(\mathbf{Y}_{-(i,t),-(j,s)}^{NT}) \omega_{js}(\mathbf{Y}_{-(i,t),-(j,s)}^{NT}) - \omega_{it}(\mathbf{Y}^{NT}) \omega_{js}(\mathbf{Y}^{NT}) \right| \middle| \mathbf{X}^{NT}, \mathbf{A}^N, \mathbf{B}^T \right] = O_P(\xi_{NT}^2).$$

REMARK 4.1. (Assumption 4.1) Part (i) guarantees that $X_{it} \neq x$ and $n_{it} = 0$ only happens for a small fraction of observations (i, t) . We are therefore able to construct proper counterfactuals $\widetilde{Y}_{it}(x)$ for most observations. Part (ii) is a boundedness condition that is standard in the matrix completion literature. Part (iii) is an independence condition that is convenient to simplify the derivations but can be generalized to weak correlation across both i and t . We use part (iv) to bound the error terms of the Taylor expansions for the bias. Part (v) imposes an injectivity condition. The functions $\mathbf{a} \mapsto \boldsymbol{\lambda}(\mathbf{a})$ and $\mathbf{b} \mapsto \mathbf{f}(\mathbf{b})$ need to be such that \mathbf{A}_i and \mathbf{B}_t can be uniquely recovered from $\boldsymbol{\lambda}_i = \boldsymbol{\lambda}(\mathbf{A}_i)$ and $\mathbf{f}_t = \mathbf{f}(\mathbf{B}_t)$. A necessary condition is that the dimensions of $\boldsymbol{\lambda}_i$ and \mathbf{f}_t are greater than or equal to the dimensions of \mathbf{A}_i and \mathbf{B}_t , respectively. This holds in our factor structure approximation when let R grow with the sample size, provided that the dimensions of \mathbf{A}_i and \mathbf{B}_t are fixed. Part (vi) holds if $\widehat{\boldsymbol{\lambda}}_i - \boldsymbol{\lambda}_i$ and $\widehat{\mathbf{f}}_t - \mathbf{f}_t$ are of order $N^{-1/2}$ and $T^{-1/2}$. We expect this assumption to be satisfied for rates $\xi_{NT} \gg \max(N^{-1}, T^{-1})$. The bandwidth parameters τ_{NT} and ν_{NT} should not be chosen too large according to part (vii). For example, if we want to achieve a rate $\xi_{NT} \ll \max(N^{-1/2}, T^{-1/2})$, then we need $\tau_{NT} \ll \max(N^{-1/4}, T^{-1/4})$ and $\nu_{NT} \ll \max(N^{-1/4}, T^{-1/4})$. Part (viii) requires that any given outcome Y_{it} is not chosen too often with too high weight in the construction of the counterfactuals $\widetilde{Y}_{js}(x)$. Finally, part (ix) is a high-level assumption that could be justified by appropriate distributional assumptions on X_{it} , \mathbf{A}_i , \mathbf{B}_t , and on the estimators $\widehat{\boldsymbol{\lambda}}_i$ and $\widehat{\mathbf{f}}_t$. We prefer to present it as a high-level assumption, because formally working out the distributional assumptions is quite cumbersome. Intuitively, if n_{it} is sufficiently large, then changing \mathbf{Y}^{NT} to $\mathbf{Y}_{-(i,t),-(j,s)}^{NT}$ should not change the constructions of the counterfactual $\widehat{Y}_{it}(x)$ very much. If that is true for all (i, t) , then the weights $\omega_{it}(\mathbf{Y}^{NT})$ should be very close to the weights $\omega_{it}(\mathbf{Y}_{-(i,t),-(j,s)}^{NT})$ and the assumption is satisfied.

THEOREM 4.1. Under Assumptions 2.1 and 4.1,

$$\widetilde{\mu}(x) - \mu(x) = O_P(\xi_{NT}).$$

As discussed in the above remark, one can achieve rates $\xi_{NT} \ll \max(N^{-1/2}, T^{-1/2})$ for sufficiently regular data generating processes, and if the bandwidth parameters τ_{NT} and ν_{NT} are chosen sufficiently small. By contrast, the low-rank approximation bias in $\widehat{\mu}(x)$ will usually prevent us

from achieving such a convergence rate for $\widehat{\mu}(x)$. This finding is consistent with our Monte Carlo results in Section 5.2, where $\widehat{\mu}(x)$ is found to typically have much smaller bias than $\widehat{\mu}(x)$.

5. NUMERICAL EXAMPLES

5.1. Election day registration and voter turnout

We illustrate the methods of the paper with an empirical application to the effect of allowing voter registration during the election day on voter turnout in the US (Xu, 2017). Voting in the US used to require registration prior to the election day in most states. Registration increased the cost of voting and was considered as one possible reason for low turnout rates. In response, some states implemented EDR laws that allowed eligible voters to register on election day when they arrived at the polling stations. These laws were not passed by all the states, and there was variation in the time of adoption across states. Thus, they were enacted by Maine, Minnesota, and Wisconsin in 1976; Wyoming, Indiana, and New Hampshire in 1994, and Connecticut in 2012.

We use a dataset on the 24 presidential elections for 47 states between 1920 and 2012 collected by Xu (2017). It includes state-level information about the turnout rate, Y_{it} , measured as the total ballots counted divided by voting-age population in state i at election t , and a treatment indicator for EDR, X_{it} , that equals one if the state i has an EDR law enacted at election t . Following Xu (2017), we exclude North Dakota where registration was never needed, and Alaska and Hawaii that were not states until 1959. Since there are only nine states that are ever treated and the treatment started in the 1976 election, we focus on effects on the treated at the elections between 1976 and 2012. We estimate average treatment effects and quantile treatment effects at multiple quantile indices.

Figure 1 compares the average turnout of states that are ever treated with states that are never treated in elections prior to the first implementation of the EDR laws in 1976. It shows that ever treated states have higher turnout rates on average than never treated states without the EDR treatment. We consider several methods to deal with this likely nonrandom assignment of EDR to estimate the ATTs (average treatment effects on the treated) for each election after 1976. First, we do a naive comparison of means between treated and nontreated states in each election (Dmeans). Second, we consider a difference-in-differences method that uses the nontreated states as controls at each election (DiD). In particular, we estimate the effects from a linear regression with state effects and election effects interacted with a EDR indicator. This method yields the ATT for each election under a parallel trend assumption between treated and nontreated states.⁵ Third, we compute our estimator based on matrix completion methods without debiasing (MC) with additive state and election effects and the parameter ρ such that the number of factors is $R = 6$. Fourth, we debias the MC estimates using the two-way matching method with 10 matches (TWM-10). Fifth, we consider the simple matching method with 5 matches (SM-5). We choose the number of matches roughly based on the numerical simulations of Section 5.2.

Figure 2 reports the estimates of the ATT of EDR at each election. The methods that account for possible nonrandom assignment of the EDR produce lower estimates of the effect than the naive comparison of means between treated and nontreated states. This finding agrees with the pre-EDR differences found in Figure 1. MC, TWM-10, and SM-5 estimates are generally larger and more stable across elections than DiD estimates. According to TWM-10, EDR laws increase

⁵ The DiD model is a special case of our model with additive effects. In this case, it imposes that there are only additive state and election effects that are the same for both treatment levels.




Figure 1. Pre-trends in turnout rate.




Figure 2. Average treatment effect on treated.




Figure 3. Time-averaged QTT.

voter turnout between 5 and 9% depending on the election. This effect is an economically significant relative to 55%, the average turnout rate for states without EDR. The estimates of the election-aggregated ATTs are 10.71%, 0.67%, 7.35%, 5.56%, and 4.87% for Dmeans, DiD, MC, TWM-10, and SM-3, respectively.

Figure 3 plots the estimates of the election-aggregated quantile treatment effect on the treated (QTT) of EDR as a function of the quantile index. We report estimates from four methods: a naive comparison of quantiles between treated and nontreated states (Dquantiles); our estimator based on matrix completion methods without debiasing (MC) with additive state and election effects and the parameter ρ such that the number of factors is $R = 3$; two-way matching with 10 matches (TWM-10); and simple matching with 5 matches (SM-5). The QTT is the difference of the quantiles between the observed turnout for the treated observations and the corresponding potential turnout had they not been treated. The quantiles of the observed turnout are estimated using sample quantiles. The estimates of the quantiles of the potential outcomes are obtained by inverting the corresponding estimates of the distribution, which are obtained by our methods replacing Y_{it} by the indicator $\mathbb{1}(Y_{it} \leq y)$ and repeating the procedure over a grid of values of y that includes the sample quantiles of observed turnout with indices $\{0.10, 0.11, \dots, 0.98\}$.⁶ Here, we find that the effect of EDR is decreasing across the distribution of turnout and ranges between 10 and 0% according to TWM-10. EDR is therefore more effective at the bottom of the voter turnout distribution. Comparing with the Dquantiles estimates, we find that the sign of the selection bias switches from positive to negative around the middle of the turnout distribution.

⁶ We rearrange the estimates of the distribution to guarantee that they are increasing with respect to y (Chernozhukov et al., 2010).

Table 1. Results for $\mu(0 | \{1\})$.

	Bias	St. Dev.	RMSE
Dmeans	0.59	0.02	0.59
DiD	0.70	0.03	0.70
MC	0.74	0.02	0.74
TWM-1	0.03	0.14	0.14
TWM-5	0.03	0.11	0.12
TWM-10	0.04	0.10	0.11
TWM-30	0.07	0.09	0.12
SM-1	0.12	0.10	0.16
SM-5	0.15	0.07	0.17
SM-10	0.19	0.06	0.20
SM-30	0.31	0.05	0.31

Notes: based on 1,000 simulations.

5.2. Monte Carlo simulations

To evaluate the performance of our methods in a controlled synthetic environment, we generate potential outcomes from an additive linear model where

$$Y_{it}(x) = x + g(A_i, B_t) + U_{it}(x), \quad x \in \{0, 1\}, i \in \{1, \dots, 30\}, t \in \{1, \dots, 30\},$$

$U_{it}(x) \sim N(0, 1/4)$ independently over i, t and x , $A_i \sim U(0, 1)$ independently over i , $B_t \sim U(0, 1)$ independently over t , $U_{it}(x)$, A_j and B_s are independent for all i, t, j and s , and g is the Gaussian kernel, i.e.,

$$g(a, b) = \frac{1}{\sqrt{2\pi}\sigma} \exp\left(-\frac{(a-b)^2}{\sigma^2}\right).$$

This design is similar to that used in Bordenave et al. (2020), with kernel function specification from the numerical simulations in Griebel and Harbrecht (2013).⁷ The parameter σ controls the decay of the singular values of g and can be calibrated to make sure the singular values decay slowly. Smaller values for σ lead to greater dispersion in the kernel function $(a, b) \mapsto g(a, b)$ and a slower singular value decay, hence can be interpreted as a measure of smoothness.⁸ The assignment of X_{it} that determines what potential outcomes are observed is similar to the election application. In particular, only observations for the first half of the units, $i \in \{1, \dots, 15\}$, and the second half of the panel, $t \in \{15, \dots, 30\}$, may be treated. For these observations, X_{it} is related to the unobserved effects (A_i, B_t) via $X_{it} = \mathbb{1}\{g(A_i, B_t) \geq c\}$, where c is a constant calibrated to $\Pr(g(A_i, B_t) \geq c) = 0.5$.

We apply similar methods to Section 5.1 to estimate the CASFs $\mu_t(0 | \{1\})$, $t \in \{15, \dots, 30\}$, and $\mu(0 | \{1\})$ using the observed variables X_{it} and $Y_{it} = Y_{it}(X_{it})$. Thus, we consider Dmeans, DiD, MC without additive effects and with the parameter ρ such that $R = 5$, and multiple versions of TWM and SM with the number of matches equal to 1, 5, 10, and 30. For each method, we compute the bias, standard deviation and RMSE (root mean square error) from 1,000 simulations. Across the simulations, we redraw the values of $U_{it}(x)$ and hold A_i , B_t and X_{it} fixed. Table 1

⁷ We find similar results in a multiplicative model where $Y_{it}(x) = (1+x)g(A_i, B_t) + U_{it}(x)$. We omit these results for the sake of brevity.

⁸ Smoothness here is specifically related to numerical smoothness, i.e., variability in the function within close neighbourhoods of its arguments.




Figure 4. Results for $t \mapsto \mu_t(0 | \{1\})$.

reports the results for the time-aggregated CASF, $\mu(0 | \{1\})$, and Figure 4 plots the results for the CASF, $\mu_t(0 | \{1\})$, as a function of t . The results show that Dmeans, DiD and MC are severely biased relative to their standard deviations. All the matching estimators reduce bias and rmse, despite increasing dispersion. As one would expect, increasing the number of matches reduces the variability of the matching estimators but increases their biases. The number of matches that minimizes the rmse is larger for the TWM than for the SM. Overall, these small-sample findings agree with the asymptotic results of Sections 3.3 and 4.

ACKNOWLEDGEMENTS

This paper was prepared for *The Econometrics Journal* Special Session on 'Econometrics of Panel Data' at the Royal Economic Society 2019 Annual Conference in Warwick University. We thank the editor Jaap Abbring, two anonymous referees, Shuowen Chen, and the participants of this conference and the 25th International Panel Data Conference for comments. This research was supported by the Economic and Social Research Council through the ESRC Centre for Microdata Methods and Practice grant RES-589-28-0001, and by the European Research Council grants ERC-2014-CoG-646917-ROMIA and ERC-2018-CoG-819086-PANEDA.

REFERENCES

- Amjad, M., D. Shah and D. Shen (2018). Robust synthetic control. *Journal of Machine Learning Research* 19(22), 1–51.
- Athey, S., M. Bayati, N. Doudchenko, G. Imbens and K. Khosravi (2021). Matrix completion methods for causal panel data models. *Journal of the American Statistical Association*, 1–41.
- Auerbach, E. (2019). Identification and estimation of a partially linear regression model using network data. arXiv eprint [arXiv:1903.09679](https://arxiv.org/abs/1903.09679).

- Bai, J. and S. Ng (2019a). Matrix completion, counterfactuals, and factor analysis of missing data. arXiv eprint [arXiv:1910.06677](https://arxiv.org/abs/1910.06677).
- Bai, J. and S. Ng (2019b). Rank regularized estimation of approximate factor models. *Journal of Econometrics* 212(1), 78–96.
- Bai, Z. D., J. W. Silverstein and Y. Q. Yin (1988). A note on the largest eigenvalue of a large dimensional sample covariance matrix. *Journal of Multivariate Analysis* 26(2), 166–68.
- Beyhum, J. and E. Gautier (2019). Square-root nuclear norm penalized estimator for panel data models with approximately low-rank unobserved heterogeneity. arXiv eprint [arXiv:1904.09192](https://arxiv.org/abs/1904.09192).
- Bordenave, C., S. Coste and R. R. Nadakuditi (2020). Detection thresholds in very sparse matrix completion. arXiv eprint [arXiv:2005.06062](https://arxiv.org/abs/2005.06062).
- Cai, J.-F., E. J. Candès and Z. Shen (2010). A singular value thresholding algorithm for matrix completion. *SIAM Journal on Optimization* 20(4), 1956–82.
- Candès, E. J. and B. Recht (2009). Exact matrix completion via convex optimization. *Foundations of Computational Mathematics* 9(6), 717.
- Candès, E. J. and T. Tao (2010). The power of convex relaxation: near-optimal matrix completion. *IEEE Transactions on Information Theory* 56(5), 2053–80.
- Chamberlain, G. (1982). Multivariate regression models for panel data. *Journal of Econometrics* 18(1), 5–46.
- Chan, M. K. and S. Kwok (2021). The PCDID approach: difference-in-differences when trends are potentially unparallel and stochastic. *Journal of Business & Economic Statistics*, 1–43.
- Chatterjee, S. (2015). Matrix estimation by universal singular value thresholding. *Annals of Statistics* 43(1), 177–214.
- Chen, M., I. Fernández-Val and M. Weidner (2021). Nonlinear factor models for network and panel data. *Journal of Econometrics* 220(2), 296–324.
- Chernozhukov, V., I. Fernández-Val and A. Galichon (2010). Quantile and probability curves without crossing. *Econometrica* 78(3), 1093–125.
- Chernozhukov, V., I. Fernández-Val, J. Hahn and W. Newey (2013). Average and quantile effects in non-separable panel models. *Econometrica* 81(2), 535–80.
- Chernozhukov, V., C. Hansen, Y. Liao and Y. Zhu (2018). Inference for heterogeneous effects using low-rank estimation of factor slopes. arXiv eprint [arXiv:1812.08089](https://arxiv.org/abs/1812.08089).
- Chernozhukov, V., C. Hansen, Y. Liao and Y. Zhu (2020). *Inference for Low-Rank Models*. Department of Economics and Center for Statistics and Data Science, MIT.
- Dzemeski, A. (2019). An empirical model of dyadic link formation in a network with unobserved heterogeneity. *Review of Economics and Statistics* 101(5), 763–76.
- Evdokimov, K. (2010). *Identification and Estimation of a Nonparametric Panel Data Model with Unobserved Heterogeneity*. Department of Economics, Princeton University.
- Fazel, S. M. (2003). *Matrix Rank Minimization with Applications*. Electrical Engineering Department, Stanford University.
- Freyberger, J. (2017, 09). Non-parametric panel data models with interactive fixed effects. *The Review of Economic Studies* 85(3), 1824–51.
- Gao, C., Y. Lu, H. H. Zhou, et al (2015). Rate-optimal graphon estimation. *The Annals of Statistics* 43(6), 2624–52.
- Geman, S. (1980). A limit theorem for the norm of random matrices. *Annals of Probability* 8(2), 252–61.
- Gobillon, L. and T. Magnac (2016). Regional policy evaluation: Interactive fixed effects and synthetic controls. *The Review of Economics and Statistics* 98(3), 535–51.
- Graham, B. S. (2017). An econometric model of network formation with degree heterogeneity. *Econometrica* 85(4), 1033–63.

- Graham, B. S. and J. L. Powell (2012). Identification and estimation of average partial effects in “irregular” correlated random coefficient panel data models. *Econometrica* 80(5), 2105–52.
- Griebel, M. and H. Harbrecht (2013). Approximation of bi-variate functions: singular value decomposition versus sparse grids. *IMA Journal of Numerical Analysis* 34(1), 28–54.
- Hoderlein, S. and H. White (2012). Nonparametric identification in nonseparable panel data models with generalized fixed effects. *Journal of Econometrics* 168(2), 300–14.
- Holland, P. W., K. B. Laskey and S. Leinhardt (1983). Stochastic blockmodels: First steps. *Social Networks* 5(2), 109–37.
- Honoré, B. (1992). Trimmed LAD and least squares estimation of truncated and censored regression models with fixed effects. *Econometrica* 60(3), 533–65.
- Hsiao, C., H. Steve Ching and S. Ki Wan (2012). A panel data approach for program evaluation: Measuring the benefits of political and economic integration of hong kong with mainland China. *Journal of Applied Econometrics* 27(5), 705–40.
- Imai, K. and I. S. Kim (2019). On the use of Two-Way Fixed Effects Regression Models for Causal Inference with Panel Data. *Political Analysis*, 1–11.
- Kim, D. and T. Oka (2014). Divorce law reforms and divorce rates in the USA: an interactive fixed-effects approach. *Journal of Applied Econometrics* 29(2), 231–45.
- Klopp, O. et al (2014). Noisy low-rank matrix completion with general sampling distribution. *Bernoulli* 20(1), 282–303.
- Latała, R. (2005). Some estimates of norms of random matrices. *Proceedings of the American Mathematical Society* 133(5), 1273–82.
- Li, K. (2018). *Inference for Factor Model Based Average Treatment Effects*. McCombs School of Business, University of Texas.
- Li, K. T. and D. R. Bell (2017). Estimation of average treatment effects with panel data: Asymptotic theory and implementation. *Journal of Econometrics* 197(1), 65–75.
- Li, Y., D. Shah, D. Song and C. L. Yu (2017). Nearest neighbors for matrix estimation interpreted as blind regression for latent variable model. *IEEE Transactions on Information Theory* 66(3), 1760–84.
- Ma, S., D. Goldfarb and L. Chen (2011). Fixed point and bregman iterative methods for matrix rank minimization. *Mathematical Programming* 128(1–2), 321–53.
- Manski, C. (1987). Semiparametric analysis of random effects linear models from binary panel data. *Econometrica* 55(2), 357–62.
- Mazumder, R., T. Hastie and R. Tibshirani (2010). Spectral regularization algorithms for learning large incomplete matrices. *Journal of Machine Learning Research* 11(80), 2287–322.
- Menzel, K. (2018). *Bootstrap with Cluster-Dependence in Two or More Dimensions*. Department of Economics, New York University.
- Moon, H. R. and M. Weidner (2017). Dynamic linear panel regression models with interactive fixed effects. *Econometric Theory* 33(1), 158–95.
- Moon, H. R. and M. Weidner (2018). Nuclear norm regularized estimation of panel regression models. arXiv eprints arXiv:1810.10987.
- Negahban, S. and M. J. Wainwright (2012). Restricted strong convexity and weighted matrix completion: optimal bounds with noise. *The Journal of Machine Learning Research* 13(1), 1665–97.
- Orbanz, P. and D. M. Roy (2015). Bayesian models of graphs, arrays and other exchangeable random structures. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 37(2), 437–61.
- Rennie, J. D. M. and N. Srebro (2005). Fast maximum margin matrix factorization for collaborative prediction. In *Proceedings of the 22nd International Conference on Machine Learning, ICML-05*, pp. 713–9, New York, NY, USA: Association for Computing Machinery.

- Silverstein, J. W. (1989). On the eigenvectors of large dimensional sample covariance matrices. *Journal of Multivariate Analysis* 30(1), 1–16.
- Srebro, N. and T. Jaakkola (2003). Weighted low-rank approximations. In *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, 720–7, Menlo Park, CA, USA: The AAAI Press.
- Wolfe, P. J. and S. C. Olhede (2013). Nonparametric graphon estimation. arXiv eprint arXiv:1309.5936.
- Xiong, R. and M. Pelger (2019). Large dimensional latent factor modeling with missing observations and applications to causal inference. arXiv eprint arXiv:1910.08273.
- Xu, J., L. Massouli and M. Lelarge (2014). Edge label inference in generalized stochastic block models: from spectral theory to impossibility results. arXiv eprint arXiv:1406.6897.
- Xu, Y. (2017). Generalized synthetic control method: Causal inference with interactive fixed effects models. *Political Analysis* 25(1), 57–76.
- Yin, Y. Q., Z. D. Bai and P. Krishnaiah (1988). On the limit of the largest eigenvalue of the large-dimensional sample covariance matrix. *Probability Theory Related Fields* 78, 509–21.
- Zelenev, A. (2020). *Identification and Estimation of Network Models with Nonparametric Unobserved Heterogeneity*. Department of Economics, Princeton University.

SUPPORTING INFORMATION

Additional Supporting Information may be found in the online version of this article at the publisher's website:

Replication Package

Managing editor Jaap Abbring handled this manuscript.

APPENDIX A: PROOFS OF RESULTS

We start with a preliminary result that relates the nuclear norm of $\Gamma^\infty(x)$ with the sum of the singular values of the function $(a, b) \mapsto m(x, a, b)$. This link will be useful to bound the approximation error of $\widehat{\Gamma}(x)$. We define

$$\|m(x, \cdot, \cdot)\|_* := \sum_{j=1}^{\infty} s_j(x).$$

LEMMA A.1. *Let Assumptions 3.1 and 3.2 hold. Then, as $N, T \rightarrow \infty$,*

$$\|\Gamma^\infty(x)\|_1 \leq \sqrt{NT} \|m(x, \cdot, \cdot)\|_* + o_P(\sqrt{NT}) = O_P(\sqrt{NT}).$$

Lemma A.1 implies that $\|\Gamma^\infty(x)\|_1$ grows with N and T at the same rate as any low-rank matrix M with elements that are of order one with bounded second moments such that $\|M\|_1 \leq \sqrt{\text{rank}(M)} \|M\|_2 = \sqrt{\text{rank}(M) \sum_{i=1}^N \sum_{t=1}^T M_{it}^2} = O_P(\sqrt{NT})$. This result will be useful for the proofs of Lemma 3.1 and of Theorem 3.1. The proof of Lemma A.1 is provided at the end of the Appendix.

The following technical lemma provides the key step in the proof of Lemma 3.1 in the main text.

LEMMA A.2. *Under Assumptions 3.1 and 3.2,*

$$\frac{1}{n(x)} \sum_{(i,t) \in \mathbb{D}(x)} (\widehat{\Gamma}_{it}(x) - \Gamma_{it}^\infty(x))^2 \leq \frac{2\rho \|\Gamma^\infty(x)\|_1}{n(x)} - \frac{2}{n(x)} \sum_{(i,t) \in \mathbb{D}(x)} \Gamma_{it}^\infty(x) E_{it},$$

for all $\rho \geq \|E(x)\|_\infty$.

Notice that Lemma A.2 is a nonstochastic finite-sample result, which only requires that $E_{it}(x)$ and $\widehat{\Gamma}(x)$ are as defined in (3.7) and (3.8). The proof of Lemma A.2 is provided at the end of the Appendix. We are now ready to provide the proof of the lemma in the main text.

Proof of Lemma 3.1. The definition of $E_{it}(x)$ in (3.7) guarantees that $E[E_{it}(x) | A^N, B^T, X^{NT}] = 0$, and Assumption 3.3 furthermore guarantees that $E_{it}(x)$ is independent across i and t and has a finite fourth moment, conditional on X^{NT}, A^N and B^T . Furthermore, $\Gamma_{it}^\infty(x) = m(x, A_i, B_t)$ only depends on A^N and B^T . We therefore find

$$\begin{aligned} & E \left[\left(\frac{1}{n(x)} \sum_{(i,t) \in \mathbb{D}(x)} \Gamma_{it}^\infty(x) E_{it} \right)^2 \middle| A^N, B^T, X^{NT} \right] \\ &= \frac{1}{n^2(x)} \sum_{(i,t) \in \mathbb{D}(x)} [\Gamma_{it}^\infty(x)]^2 E[E_{it}^2 | A^N, B^T, X^{NT}] \\ &\leq \frac{b^{1/2}}{n^2(x)} \sum_{(i,t) \in \mathbb{D}(x)} [\Gamma_{it}^\infty(x)]^2 = O_P(1/n(x)), \end{aligned}$$

where b is the constant from Assumption 3.3. From this we conclude that

$$\frac{1}{n(x)} \sum_{(i,t) \in \mathbb{D}(x)} \Gamma_{it}^\infty(x) E_{it} = O_P\left(\frac{1}{n^{1/2}(x)}\right) = o_P(1). \tag{A.1}$$

Next, applying Assumption 3.3 and theorem 2 in Latała (2005) we find

$$\begin{aligned} E[\|E(x)\|_\infty \mid A^N, B^T, X^{NT}] &\leq C \left\{ \max_t \sqrt{\sum_i E[E_{it}(x)^2 \mid A^N, B^T, X^{NT}]} \right. \\ &\quad + \max_i \sqrt{\sum_t E[E_{it}(x)^2 \mid A^N, B^T, X^{NT}]} \\ &\quad \left. + \left(\sum_{i,t} E[E_{it}(x)^4 \mid A^N, B^T, X^{NT}] \right)^{1/4} \right\} \\ &\leq C b^{1/4} \left\{ \sqrt{N} + \sqrt{T} + n(x)^{1/4} \right\} = O_p(\sqrt{N+T}), \end{aligned}$$

where C is a universal constant. We therefore have $\|E(x)\|_\infty = O_p(\sqrt{N+T})$, and since we assume that $\rho = \rho_{NT}$ satisfies $\rho_{NT}/\sqrt{N+T} \rightarrow \infty$ we conclude that

$$\rho_{NT} \geq \|E(x)\|_\infty$$

with probability approaching one. We can therefore apply Lemma A.2 to find that, with probability approaching one, we have

$$\begin{aligned} \frac{1}{n(x)} \sum_{(i,t) \in \mathbb{D}(x)} (\widehat{\Gamma}_{it}(x) - \Gamma_{it}^\infty(x))^2 &\leq \frac{2\rho_{NT} \|\Gamma^\infty(x)\|_1}{n(x)} - \frac{2}{n(x)} \sum_{(i,t) \in \mathbb{D}(x)} \Gamma_{it}^\infty(x) E_{it} \\ &= \frac{2\rho_{NT} O_p(\sqrt{NT})}{n(x)} + o_p(1) \\ &= o_p(1), \end{aligned}$$

where we applied (A.1) and Lemma A.1, as well as the condition $\rho_{NT}\sqrt{NT}/n(x) \rightarrow 0$. □

In the following consider a generic reduced-form parameter

$$v_0(x) = \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T W_{it}(x) \Gamma_{it}^\infty(x), \tag{A.2}$$

with corresponding estimator

$$\widehat{v}(x) = \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T W_{it}(x) \widehat{\Gamma}_{it}(x), \tag{A.3}$$

where $W_{it}(x)$ are given weights.

The following proposition provides a finite-sample nonstochastic bound for the error of this reduced-form estimator.

PROPOSITION A.1. *Let the Assumptions 3.1, 3.2, and 3.3 hold. Let $P_{it}(x)$ be nonzero real numbers for all $(i, t) \in \mathbb{N} \times \mathbb{T}$. Define*

$$\begin{aligned}
 V_{it}(x) &:= \frac{W_{it}(x) P_{it}^{-1}(x)(D_{it}(x) - P_{it}(x))}{\frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T W_{it}(x)^2 P_{it}^{-1}(x)}, \\
 c_1 &:= \frac{1 - \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T W_{it}(x) P_{it}^{-1}(x) V_{it}(x)}{\frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T W_{it}(x)^2 P_{it}^{-1}(x)}, \\
 c_2 &:= \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T V_{it}(x) \Gamma_{it}^\infty(x), \\
 c_3 &:= \frac{2\rho}{c_1 NT} \|\Gamma^\infty(x)\|_1 - \frac{2}{c_1 NT} \sum_{(i,t) \in \mathbb{D}(x)} E_{it}(x) \Gamma_{it}^\infty(x) + \left(\frac{c_2}{c_1}\right)^2, \\
 c_4 &:= \sqrt{c_3} + \frac{|c_2|}{c_1},
 \end{aligned}$$

and let $V(x)$ be the $N \times T$ matrix with elements $V_{it}(x)$. If $c_1 > 0$ and $\rho > \|\mathbf{E}(x)\|_\infty + c_4 \|V(x)\|_\infty$, then

$$|\widehat{v}(x) - v_0(x)| \leq c_4.$$

The proof of Proposition A.1 is provided at the end of the Appendix. Proposition A.1 is the key step required for the proof of Theorem 3.1. However, before proving this main text result we want to provide an informal remark on the usefulness of Proposition A.1 more generally.

REMARK 5.1. (Consistency of $\widehat{v}(x)$) Proposition A.1 holds for all $P_{it}(x) \in \mathbb{R} \setminus \{0\}$, but for the proposition to be useful in showing consistency of $\widehat{v}(x)$ we need to choose $P_{it}(x)$ such that c_2 and $\|V(x)\|_\infty$ are not too large. The easiest way to guarantee this is to consider X_{it} to be random and weakly correlated across both i and t , and to define $P_{it}(x)$ as the propensity score; that is,

$$P_{it}(x) = \Pr(X_{it} = x \mid \mathbf{A}^N, \mathbf{B}^T),$$

which is assumed to be positive and not too small—e.g., we need that

$$q := \left[\frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T W_{it}(x)^2 P_{it}^{-1}(x) \right]^{-1},$$

converges to some positive constant. Then $V_{it}(x)$ has mean zero, analogous to $E_{it}(x)$, and

$$\begin{aligned}
 c_1 &= q + O_P(1/\sqrt{NT}), \\
 c_2 &= O_P(1/\sqrt{NT}) \\
 c_3 &= \frac{2\rho}{q NT} \|\Gamma^\infty(x)\|_1 + O_P(1/\sqrt{NT}), \\
 c_4 &= \sqrt{\frac{2\rho}{q NT} \|\Gamma^\infty(x)\|_1} + \text{smaller order terms.}
 \end{aligned}$$

Thus, if, as in Lemma 3.1, $\rho = \rho_{NT}$ such that $\rho_{NT}/\sqrt{N+T} \rightarrow \infty$ and $\rho_{NT}/\sqrt{NT} \rightarrow 0$ as $N, T \rightarrow \infty$, then

$$\widehat{v}(x) = v_0(x) + o_P(1).$$

The following proof formalizes this heuristic argument for the case that $W_{it}(x) = 1$.

Proof of Theorem 3.1. Let $W_{it}(x) = 1$, and let $v_0(x)$ and $\widehat{v}(x)$ be as defined in (A.2) and (A.3) above. We then have

$$\begin{aligned} \mu(x) &= v_0(x), \\ \widehat{\mu}(x) &= \widehat{v}(x) + \frac{1}{NT} \sum_{(i,t) \in \mathbb{D}(x)} E_{it}(x) - \frac{1}{NT} \sum_{(i,t) \in \mathbb{D}(x)} [\widehat{\Gamma}_{it}(x) - \Gamma_{it}^\infty(x)]. \end{aligned} \tag{A.4}$$

We drop all the arguments x in the rest of this proof. We want to apply Proposition A.1 with $P_{it} = \Pr(X_{it} = x \mid A^N, \mathbf{B}^T) > 0$. Let $G_{it} = P_{it}^{-1}(D_{it} - P_{it})$ be as defined in Theorem 3.1, and also define $q := \left[\frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T P_{it}^{-1} \right]^{-1}$. Since $P_{it} \in [0, 1]$ we also have $q \in [0, 1]$, and the theorem assumes that $q^{-1} = O_P(1)$. Using Lemma A.1 we know that $\|\Gamma^\infty\|_1 = O_P(\sqrt{NT})$, and we have already found that $\sum_{(i,t) \in \mathbb{D}} \Gamma_{it}^\infty E_{it} = O_P(n^{1/2})$ in (A.1) above. Using this together the other assumptions in the theorem we find that

$$\begin{aligned} V_{it} &= q G_{it} \\ c_1 &= q \left(1 - \frac{q}{NT} \sum_{i=1}^N \sum_{t=1}^T P_{it}^{-1} G_{it} \right) = q [1 - o_P(1)], \\ c_2 &= \frac{q}{NT} \sum_{i=1}^N \sum_{t=1}^T G_{it} \Gamma_{it}^\infty = o_P(1), \\ c_3 &= \frac{2\rho O_P(\sqrt{NT})}{c_1 NT} - \frac{O_P(n^{1/2})}{c_1 NT} + \left(\frac{c_2}{c_1} \right)^2 = o_P(1), \\ c_4 &= \sqrt{c_3} + \frac{|c_2|}{c_1} = o_P(1). \end{aligned}$$

We furthermore have

$$\|\mathbf{V}\|_\infty = q \|\mathbf{G}\|_\infty = O_P(1) O_P(\sqrt{N+T}) = O_P(\sqrt{N+T}).$$

In the proof of Lemma 3.1 we already argued that $\|\mathbf{E}\|_\infty = O_P(\sqrt{N+T})$. Since we assume that $\rho = \rho_{NT}$ satisfies $\rho_{NT}/\sqrt{N+T} \rightarrow \infty$ we conclude that

$$\rho > \|\mathbf{E}\|_\infty + c_4 \|\mathbf{V}\|_\infty,$$

with probability approach one. We can therefore apply Proposition A.1 to find that with probability approach one we have

$$|\widehat{v} - v_0| \leq c_4 = o_P(1).$$

We have thus shown that $\widehat{v} = v_0 + o_P(1)$.

Furthermore, analogous to the result in (A.1) we can show that $\sum_{(i,t) \in \mathbb{D}} E_{it} = O_P(n^{1/2})$, and we therefore have $\frac{1}{NT} \sum_{(i,t) \in \mathbb{D}} E_{it} = o_P(1)$. Finally, applying Lemma 3.1 we have Next, from which we know that

$$\left[\frac{1}{n} \sum_{(i,t) \in \mathbb{D}} (\widehat{\Gamma}_{it} - \Gamma_{it}^\infty) \right]^2 \leq \frac{1}{n} \sum_{(i,t) \in \mathbb{D}} (\widehat{\Gamma}_{it} - \Gamma_{it}^\infty)^2 = o_P(1),$$

and therefore $\frac{1}{NT} \sum_{(i,t) \in \mathbb{D}(x)} [\widehat{\Gamma}_{it}(x) - \Gamma_{it}^\infty(x)] = o_P(1)$. Plugging those results into (A.4) we find $\widehat{\mu}(x) = \mu(x) + o_P(1)$. \square

In this section we present and prove a more general version of Theorem 4.1. Let $\phi_i = \phi(x, A_i)$ and $\psi_t = \psi(x, B_t)$ be transformations of A_i and B_t . Let $\widehat{\phi}_i$ and $\widehat{\psi}_t$ be corresponding estimators. In the main

text we presented the special case where $\widehat{\phi}_i$ and $\widehat{\psi}_t$ were equal to the factor loadings and factors obtained from $\widehat{\Gamma}(x)$, but many other choices of $\widehat{\phi}_i$ and $\widehat{\psi}_t$ are conceivable. We again define

$$\mathbb{N}_i = \{j \in \mathbb{N} \setminus \{i\} : \|\widehat{\phi}_i - \widehat{\phi}_j\| \leq \tau_{NT}\}, \quad \mathbb{T}_t = \{s \in \mathbb{T} \setminus \{t\} : \|\widehat{\psi}_t - \widehat{\psi}_s\| \leq \nu_{NT}\},$$

for some bandwidth parameters $\tau_{NT} > 0$ and $\nu_{NT} > 0$. A debiased estimator of the reduced-form parameter in (A.2) is given by

$$\widetilde{v}(x) = \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T W_{it}(x) \widetilde{Y}_{it}(x),$$

where $\widetilde{Y}_{it}(x)$ is defined as in (4.1). In the main text we only discussed the special case $W_{it}(x) = 1$. We can write $\widetilde{v}(x)$ as

$$\widetilde{v}(x) = \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \omega_{it} Y_{it},$$

where the weights ω_{it} are functions of $\widehat{\phi}_j$ and $\widehat{\psi}_s$ for all $j \in \mathbb{N}$ and $s \in \mathbb{T}$. Assumption 4.1 in the main text is generalized as follows.

ASSUMPTION A.1. *There exists a sequence $\xi_{NT} > 0$ such that $\xi_{NT} \rightarrow 0$ as $N, T \rightarrow \infty$, and*

- (i) $\frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T W_{it}(x) \mathbb{1}\{X_{it} \neq x \text{ \& } n_{it} = 0\} = O_P(\xi_{NT})$.
- (ii) Y_{it} and $W_{it}(x)$ are uniformly bounded over i, t, N, T .
- (iii) Y_{it} is independent across both i and t , conditional on $\mathbf{X}^{NT}, \mathbf{A}^N, \mathbf{B}^T$.
- (iv) The function $(\mathbf{a}, \mathbf{b}) \mapsto m(x, \mathbf{a}, \mathbf{b})$ is twice continuously differentiable with uniformly bounded second derivatives.
- (v) There exists $c > 0$ such that $\|\mathbf{a}_1 - \mathbf{a}_2\| \leq c \|\phi(\mathbf{a}_1) - \phi(\mathbf{a}_2)\|$ for all $\mathbf{a}_1, \mathbf{a}_2 \in \mathbb{A}$, and $\|\mathbf{b}_1 - \mathbf{b}_2\| \leq c \|\psi(\mathbf{b}_1) - \psi(\mathbf{b}_2)\|$ for all $\mathbf{b}_1, \mathbf{b}_2 \in \mathbb{B}$.
- (vi) $\frac{1}{N} \sum_{i=1}^N \left(\|\widehat{\phi}_i - \phi_i\|^2 + \max_{j \in \mathbb{N}_i} \|\widehat{\phi}_j - \phi_j\|^2 \right) = O_P(\xi_{NT})$.
 $\frac{1}{T} \sum_{t=1}^T \left(\|\widehat{\psi}_t - \psi_t\|^2 + \max_{s \in \mathbb{T}_t} \|\widehat{\psi}_s - \psi_s\|^2 \right) = O_P(\xi_{NT})$.
- (vii) $\tau_{NT}^2 = O_P(\xi_{NT})$ and $\nu_{NT}^2 = O_P(\xi_{NT})$.
- (viii) $\frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \mathbb{E}[\omega_{it}^2 | \mathbf{X}^{NT}, \mathbf{A}^N, \mathbf{B}^T] = O_P(NT \xi_{NT}^2)$.
- (ix) Let $\mathbf{Y}_{-(i,t),-(j,s)}^{NT}$ be the outcome matrix \mathbf{Y}^{NT} , but with Y_{it} and Y_{js} replace by zero (or some other nonrandom number), and all other outcomes unchanged. We assume

$$\begin{aligned} \frac{1}{(NT)^2} \sum_{i,j=1}^N \sum_{t,s=1}^T \mathbb{1}\{(i,t) \neq (j,s)\} \mathbb{E} \left[\omega_{it}(\mathbf{Y}_{-(i,t),-(j,s)}^{NT}) \omega_{js}(\mathbf{Y}_{-(i,t),-(j,s)}^{NT}) \right. \\ \left. - \omega_{it}(\mathbf{Y}^{NT}) \omega_{js}(\mathbf{Y}^{NT}) \middle| \mathbf{X}^{NT}, \mathbf{A}^N, \mathbf{B}^T \right] = O_P(\xi_{NT}^2). \end{aligned}$$

The generalized version of Theorem 4.1 is given in the following.

THEOREM A.1. *Under Assumptions 2.1 and A.1,*

$$\widetilde{v}(x) - v_0(x) = O_P(\xi_{NT}).$$

Proof of Theorem A.1 (containing Theorem 4.1 as a special case). Define $m_{it}(x) := m(x, \mathbf{A}_i, \mathbf{B}_t)$. We decompose

$$\widetilde{v}(x) - v_0(x) = e_0(x) + e_1(x) + e_2(x), \tag{A.5}$$

where

$$e_0(x) = \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T W_{it}(x) \mathbb{1}\{X_{it} \neq x \text{ \& } n_{it} = 0\} [m_{it}(X_{it}) - m_{it}(x)],$$

and

$$e_1(x) := \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \mathbb{1}\{X_{it} \neq x \text{ \& } n_{it} > 0\} W_{it}(x) e_{1,it}(x),$$

$$e_{1,it}(x) := \frac{\sum_{j \in \mathbb{N}_i} \sum_{s \in \mathbb{T}_t} \mathbb{1}\{X_{is} = X_{jt} = X_{js} = x\} [m_{is}(x) + m_{jt}(x) - m_{js}(x) - m_{it}(x)]}{\sum_{j \in \mathbb{N}_i} \sum_{s \in \mathbb{T}_t} \mathbb{1}\{X_{is} = X_{jt} = X_{js} = x\}},$$

and

$$e_2(x) := \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \omega_{it} E_{it}.$$

In the following we consider $e_0(x)$, $e_1(x)$, $e_2(x)$ separately.

Bound on $e_0(x)$: Assumption A.1(i) and (ii) guarantee that

$$\begin{aligned} |e_0(x)| &\leq (\max_{it} |m_{it}(X_{it}) - m_{it}(x)|) \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T W_{it}(x) \mathbb{1}\{X_{it} \neq x \text{ \& } n_{it} = 0\} \\ &= O_P(\xi_{NT}). \end{aligned} \tag{A.6}$$

Bound on $e_1(x)$: Assumption A.1(iv) guarantees that there exists a constant $b > 0$ such that

$$\begin{aligned} \left| m(x, \mathbf{a}, \mathbf{b}) - m(x, \mathbf{A}_i, \mathbf{B}_t) - (\mathbf{a} - \mathbf{A}_i)' \frac{\partial m(x, \mathbf{A}_i, \mathbf{B}_t)}{\partial \mathbf{A}_i} - (\mathbf{b} - \mathbf{B}_t)' \frac{\partial m(x, \mathbf{A}_i, \mathbf{B}_t)}{\partial \mathbf{B}_t} \right| \\ \leq b (\|\mathbf{a} - \mathbf{A}_i\|^2 + \|\mathbf{b} - \mathbf{B}_t\|^2). \end{aligned}$$

Using this we find that

$$m_{is}(x) + m_{jt}(x) - m_{js}(x) - m_{it}(x) \leq 2b (\|\mathbf{A}_i - \mathbf{A}_j\|^2 + \|\mathbf{B}_t - \mathbf{B}_s\|^2),$$

and therefore

$$\begin{aligned} |e_{1,it}(x)| &\leq \frac{2b \sum_{j \in \mathbb{N}_i} \sum_{s \in \mathbb{T}_t} \mathbb{1}\{X_{is} = X_{jt} = X_{js} = x\} (\|\mathbf{A}_i - \mathbf{A}_j\|^2 + \|\mathbf{B}_t - \mathbf{B}_s\|^2)}{\sum_{j \in \mathbb{N}_i} \sum_{s \in \mathbb{T}_t} \mathbb{1}\{X_{is} = X_{jt} = X_{js} = x\}} \\ &\leq 2b \left(\max_{j \in \mathbb{N}_i} \|\mathbf{A}_i - \mathbf{A}_j\|^2 + \max_{s \in \mathbb{T}_t} \|\mathbf{B}_t - \mathbf{B}_s\|^2 \right). \end{aligned}$$

We thus find

$$\begin{aligned} |e_1(x)| &\leq 2b \left(\max_{ij} |W_{it}(x)| \right) \left(\frac{1}{N} \sum_{i=1}^N \max_{j \in \mathbb{N}_i} \|\mathbf{A}_i - \mathbf{A}_j\|^2 + \frac{1}{T} \sum_{t=1}^T \max_{s \in \mathbb{T}_t} \|\mathbf{B}_t - \mathbf{B}_s\|^2 \right) \\ &\leq 2bc \left(\max_{ij} |W_{it}(x)| \right) \left(\frac{1}{N} \sum_{i=1}^N \max_{j \in \mathbb{N}_i} \|\boldsymbol{\phi}(\mathbf{A}_i) - \boldsymbol{\phi}(\mathbf{A}_j)\|^2 \right. \\ &\quad \left. + \frac{1}{T} \sum_{t=1}^T \max_{s \in \mathbb{T}_t} \|\boldsymbol{\psi}(\mathbf{B}_t) - \boldsymbol{\psi}(\mathbf{B}_s)\|^2 \right) \\ &= 2bc \left(\max_{ij} |W_{it}(x)| \right) \left(\frac{1}{N} \sum_{i=1}^N \max_{j \in \mathbb{N}_i} \|\boldsymbol{\phi}_i - \boldsymbol{\phi}_j\|^2 + \frac{1}{T} \sum_{t=1}^T \max_{s \in \mathbb{T}_t} \|\boldsymbol{\psi}_t - \boldsymbol{\psi}_s\|^2 \right). \end{aligned}$$

Using the triangle inequality, the definition of \mathbb{N}_i , and the general inequality $(x_1 + x_2 + x_3)^2 \leq 3(x_1^2 + x_2^2 + x_3^2)$, for $x_1, x_2, x_3 \in \mathbb{R}$, we have

$$\begin{aligned} \max_{j \in \mathbb{N}_i} \|\phi_i - \phi_j\|^2 &\leq \max_{j \in \mathbb{N}_i} (\|\widehat{\phi}_i - \widehat{\phi}_j\| + \|\widehat{\phi}_i - \phi_i\| + \|\widehat{\phi}_j - \phi_j\|)^2 \\ &\leq \max_{j \in \mathbb{N}_i} (\tau_{NT} + \|\widehat{\phi}_i - \phi_i\| + \|\widehat{\phi}_j - \phi_j\|)^2 \\ &\leq 3\tau_{NT}^2 + 3\|\widehat{\phi}_i - \phi_i\|^2 + 3 \max_{j \in \mathbb{N}_i} \|\widehat{\phi}_j - \phi_j\|^2. \end{aligned}$$

Analogously we find

$$\max_{s \in \mathbb{T}_t} \|\psi_t - \psi_s\|^2 \leq 3\nu_{NT}^2 + 3\|\widehat{\psi}_t - \psi_t\|^2 + 3 \max_{s \in \mathbb{T}_t} \|\widehat{\psi}_s - \psi_s\|^2.$$

We thus obtain

$$\begin{aligned} |e_1(x)| &\leq 6bc \left(\max_{ij} |W_{it}(x)| \right) \left\{ \tau_{NT}^2 + \nu_{NT}^2 \right. \\ &\quad + \frac{1}{N} \sum_{i=1}^N \left(\|\widehat{\phi}_i - \phi_i\|^2 + \max_{j \in \mathbb{N}_i} \|\widehat{\phi}_j - \phi_j\|^2 \right) \\ &\quad \left. + \frac{1}{T} \sum_{t=1}^T \left(\|\widehat{\psi}_t - \psi_t\|^2 + \max_{s \in \mathbb{T}_t} \|\widehat{\psi}_s - \psi_s\|^2 \right) \right\} \\ &= O_P(\xi_{NT}). \end{aligned} \tag{A.7}$$

Bound on $e_2(x)$: We have

$$[e_2(x)]^2 = \frac{1}{(NT)^2} \sum_{i,j=1}^N \sum_{t,s=1}^T \omega_{it}(\mathbf{Y}^{NT}) \omega_{js}(\mathbf{Y}^{NT}) E_{it} E_{js} = T_0 + T_1 + T_2,$$

where

$$T_0 := \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T \omega_{it}^2(\mathbf{Y}^{NT}) E_{it}^2,$$

$$\begin{aligned} T_1 &:= \frac{1}{(NT)^2} \sum_{i,j=1}^N \sum_{t,s=1}^T \mathbb{1}\{(i,t) \neq (j,s)\} \\ &\quad \times [\omega_{it}(\mathbf{Y}^{NT}) \omega_{js}(\mathbf{Y}^{NT}) - \omega_{it}(\mathbf{Y}_{-(i,t),-(j,s)}^{NT}) \omega_{js}(\mathbf{Y}_{-(i,t),-(j,s)}^{NT})] E_{it} E_{js}, \end{aligned}$$

$$T_2 := \frac{1}{(NT)^2} \sum_{i,j=1}^N \sum_{t,s=1}^T \mathbb{1}\{(i,t) \neq (j,s)\} \omega_{it}(\mathbf{Y}_{-(i,t),-(j,s)}^{NT}) \omega_{js}(\mathbf{Y}_{-(i,t),-(j,s)}^{NT}) E_{it} E_{js}.$$

We have

$$\begin{aligned} \mathbb{E}[T_0 | \mathbf{X}^{NT}, \mathbf{A}^N, \mathbf{B}^T] &\leq \left(\max_{i,t} |E_{it}| \right)^2 \frac{1}{(NT)^2} \sum_{i=1}^N \sum_{t=1}^T \mathbb{E}[\omega_{it}^2(\mathbf{Y}^{NT}) | \mathbf{X}^{NT}, \mathbf{A}^N, \mathbf{B}^T] \\ &= O_P(\xi_{NT}^2), \end{aligned}$$

and

$$\begin{aligned} & \left| \mathbb{E} \left[T_1 \mid \mathbf{X}^{NT}, \mathbf{A}^N, \mathbf{B}^T \right] \right| \\ & \leq \left(\max_{i,t} |E_{it}| \right)^2 \frac{1}{(NT)^2} \sum_{i,j=1}^N \sum_{t,s=1}^T \mathbb{1} \{ (i,t) \neq (j,s) \} \\ & \quad \times \mathbb{E} \left[\left| \omega_{it}(\mathbf{Y}^{NT}) \omega_{js}(\mathbf{Y}^{NT}) - \omega_{it}(\mathbf{Y}_{-(i,t),-(j,s)}^{NT}) \omega_{js}(\mathbf{Y}_{-(i,t),-(j,s)}^{NT}) \right| \mid \mathbf{X}^{NT}, \mathbf{A}^N, \mathbf{B}^T \right] \\ & = O_P(\xi_{NT}^2), \end{aligned}$$

where we used that Y_{it} (and thus E_{it}) is uniformly bounded, together with Assumption A.1(viii) and (ix). Next, for $(i,t) \neq (j,s)$ we have

$$\begin{aligned} & \mathbb{E} \left[\omega_{it}(\mathbf{Y}_{-(i,t),-(j,s)}^{NT}) \omega_{js}(\mathbf{Y}_{-(i,t),-(j,s)}^{NT}) E_{it} E_{js} \mid \mathbf{Y}_{-(i,t),-(j,s)}^{NT}, \mathbf{X}^{NT}, \mathbf{A}^N, \mathbf{B}^T \right] \\ & = \omega_{it}(\mathbf{Y}_{-(i,t),-(j,s)}^{NT}) \omega_{js}(\mathbf{Y}_{-(i,t),-(j,s)}^{NT}) \mathbb{E} \left[E_{it} E_{js} \mid \mathbf{Y}_{-(i,t),-(j,s)}^{NT}, \mathbf{X}^{NT}, \mathbf{A}^N, \mathbf{B}^T \right] \\ & = \omega_{it}(\mathbf{Y}_{-(i,t),-(j,s)}^{NT}) \omega_{js}(\mathbf{Y}_{-(i,t),-(j,s)}^{NT}) \\ & \quad \mathbb{E} \left[E_{it} \mid \mathbf{Y}_{-(i,t),-(j,s)}^{NT}, \mathbf{X}^{NT}, \mathbf{A}^N, \mathbf{B}^T \right] \mathbb{E} \left[E_{js} \mid \mathbf{Y}_{-(i,t),-(j,s)}^{NT}, \mathbf{X}^{NT}, \mathbf{A}^N, \mathbf{B}^T \right] \\ & = 0, \end{aligned}$$

where we used $\mathbb{E} \left[E_{it} \mid \mathbf{X}^{NT}, \mathbf{A}^N, \mathbf{B}^T \right] = 0$ together with the assumption that Y_{it} (and thus E_{it}) is independent across both i and t , conditional on $\mathbf{X}^{NT}, \mathbf{A}^N, \mathbf{B}^T$. By the law of iterated expectations the last display result also implies that for $(i,t) \neq (j,s)$ we have

$$\mathbb{E} \left[\omega_{it}(\mathbf{Y}_{-(i,t),-(j,s)}^{NT}) \omega_{js}(\mathbf{Y}_{-(i,t),-(j,s)}^{NT}) E_{it} E_{js} \mid \mathbf{X}^{NT}, \mathbf{A}^N, \mathbf{B}^T \right] = 0.$$

Using this we obtain that

$$\mathbb{E} \left[T_2 \mid \mathbf{X}^{NT}, \mathbf{A}^N, \mathbf{B}^T \right] = 0.$$

Combining those results on T_0, T_1, T_2 we obtain

$$\mathbb{E} \left\{ [e_2(x)]^2 \mid \mathbf{X}^{NT}, \mathbf{A}^N, \mathbf{B}^T \right\} = O_P(\xi_{NT}^2),$$

which implies $e_2 = O_P(\xi_{NT})$. Together with (A.5), (A.6), and (A.7) this gives the statement of the theorem. \square

Proof of Lemma A.1. Let $\mathbf{u}_j(x)$ be the N -vector with elements $u_j(x, \mathbf{A}_i)$, and let $\mathbf{v}_j(x)$ be the T -vector with elements $v_j(x, \mathbf{B}_i)$. Then we have $\mathbf{\Gamma}^\infty(x) = \sum_{j=1}^\infty s_j(x) \mathbf{u}_j(x) \mathbf{v}_j^\top(x)$, and therefore

$$\begin{aligned} \|\mathbf{\Gamma}^\infty(x)\|_1 &\leq \sum_{j=1}^\infty s_j(x) \|\mathbf{u}_j(x)\| \|\mathbf{v}_j(x)\| \\ &= \sqrt{NT} \sum_{j=1}^\infty s_j(x) \sqrt{\frac{1}{N} \sum_{i=1}^N [u_j(x, \mathbf{A}_i)]^2} \sqrt{\frac{1}{T} \sum_{i=1}^T [v_j(x, \mathbf{B}_i)]^2} \\ &\leq \sqrt{NT} \sum_{j=1}^\infty s_j(x) \left(1 + \frac{\frac{1}{N} \sum_{i=1}^N [u_j(x, \mathbf{A}_i)]^2 - 1}{2}\right) \left(1 + \frac{\frac{1}{T} \sum_{i=1}^T [v_j(x, \mathbf{B}_i)]^2 - 1}{2}\right) \\ &= \sqrt{NT} \sum_{j=1}^\infty s_j(x) + \sqrt{NT} R_{NT} \\ &= \sqrt{NT} \|m(x, \cdot, \cdot)\|_* + \sqrt{NT} R_{NT}, \end{aligned}$$

where for the second inequality we used that $\sqrt{z} \leq 1 + \frac{z-1}{2}$, for all $z \geq 0$, and we defined $R_{NT} = \frac{1}{NT} \sum_{i=1}^N \sum_{i=1}^T r_{it}$, with

$$r_{it} = \sum_{j=1}^\infty s_j(x) \left\{ \frac{[u_j(x, \mathbf{A}_i)]^2 + [v_j(x, \mathbf{B}_i)]^2}{4} + \frac{[u_j(x, \mathbf{A}_i)]^2 [v_j(x, \mathbf{B}_i)]^2}{4} - \frac{3}{4} \right\}.$$

Assumption 3.2 guarantees that $[u_j(x, \mathbf{A}_i)]^2$ and $[v_j(x, \mathbf{B}_i)]^2$ have mean equal to one, which implies that r_{it} has mean zero. Assumption 3.1 and the WLLN therefore guarantees that $R_{NT} = o_p(1)$. We have thus shown that $\|\mathbf{\Gamma}^\infty(x)\|_1 \leq \sqrt{NT} \|m(x, \cdot, \cdot)\|_* + o_p(\sqrt{NT})$, and since $\|m(x, \cdot, \cdot)\|_*$ is finite and nonrandom we also have $\|\mathbf{\Gamma}^\infty(x)\|_1 = O_p(\sqrt{NT})$. \square

Proof of Lemma A.2. The nuclear norm (or trace norm) can be defined by

$$\begin{aligned} \|\mathbf{\Gamma}\|_1 &= \max_{\{\mathbf{M} \in \mathbb{R}^{N \times T} : \|\mathbf{M}\|_\infty \leq 1\}} \underbrace{\text{Tr}(\mathbf{M}' \mathbf{\Gamma})}_{= \sum_{i=1}^N \sum_{t=1}^T M_{it} \Gamma_{it}}. \end{aligned} \tag{A.8}$$

Our assumption $\rho \geq \|\mathbf{E}(x)\|_\infty$ guarantees that a possible choice in this maximization is $\mathbf{M} = \rho^{-1} \mathbf{E}(x)$, and we therefore have

$$\rho \|\mathbf{\Gamma}\|_1 \geq \sum_{i=1}^N \sum_{t=1}^T D_{it}(x) E_{it}(x) \Gamma_{it}.$$

Using this and the model $Y_{it} = \Gamma_{it}^\infty(x) + E_{it}(x)$, for $X_{it} = x$, we find that

$$\begin{aligned} Q_{NT}(\mathbf{\Gamma}, \rho, x) &= \frac{1}{2} \sum_{i=1}^N \sum_{t=1}^T D_{it}(x) (Y_{it} - \Gamma_{it})^2 + \rho \|\mathbf{\Gamma}\|_1 \\ &\geq \frac{1}{2} \sum_{i=1}^N \sum_{t=1}^T D_{it}(x) (\Gamma_{it}^\infty(x) + E_{it}(x) - \Gamma_{it})^2 + \sum_{i=1}^N \sum_{t=1}^T D_{it}(x) E_{it}(x) \Gamma_{it} \\ &= \frac{1}{2} \sum_{i=1}^N \sum_{t=1}^T D_{it}(x) (\Gamma_{it}^\infty(x) - \Gamma_{it})^2 + \sum_{i=1}^N \sum_{t=1}^T D_{it}(x) \Gamma_{it}^\infty(x) E_{it}(x) \\ &\quad + \frac{1}{2} \sum_{i=1}^N \sum_{t=1}^T D_{it}(x) E_{it}^2(x). \end{aligned}$$

By definition we have

$$Q_{NT}(\widehat{\mathbf{\Gamma}}(x), \rho, x) \leq Q_{NT}(\mathbf{\Gamma}^\infty(x), \rho, x) = \frac{1}{2} \sum_{i=1}^N \sum_{t=1}^T D_{it}(x) E_{it}^2(x) + \rho \|\mathbf{\Gamma}^\infty(x)\|_1.$$

Combining the results in the last two displays gives the statement of the lemma. □

Proof of Proposition A.1. In this proof we drop the argument x everywhere, and we define $\theta = NTv$ and $\theta_0 = NTv_0$. Define the NT -vectors $\boldsymbol{\gamma} = \text{vec}(\mathbf{\Gamma})$, $\boldsymbol{\gamma}^\infty = \text{vec}(\mathbf{\Gamma}^\infty)$, $\mathbf{w} = \text{vec}(W_{it} : i \in \mathbb{N}, t \in \mathbb{T})$, $\mathbf{d} = \text{vec}(D_{it} : i \in \mathbb{N}, t \in \mathbb{T})$, and $\mathbf{p} = \text{vec}(P_{it} : i \in \mathbb{N}, t \in \mathbb{T})$. Then, $\text{diag}(\mathbf{p})$ is an $NT \times NT$ diagonal matrix. For $\rho > 0$ and $\theta \in \mathbb{R}$ we define

$$L_{NT}(\theta, \rho) = \min_{\{\boldsymbol{\Gamma} \in \mathbb{R}^{N \times T} : \theta = \mathbf{w}'\boldsymbol{\gamma}\}} Q_{NT}(\mathbf{\Gamma}, \rho),$$

which is the profile objective function that minimizes $Q_{NT}(\mathbf{\Gamma}, \rho)$ over almost all parameters $\mathbf{\Gamma}$, only keeping our parameter of interest fixed at $\theta = \mathbf{w}'\boldsymbol{\gamma} = \sum_{i=1}^N \sum_{t=1}^T W_{it} \Gamma_{it}$. Our goal is to show that the minimizing value

$$\widehat{\theta} := \underset{\theta \in \mathbb{R}}{\text{argmin}} L_{NT}(\theta, \rho) = \sum_{i=1}^N \sum_{t=1}^T W_{it} \widehat{\Gamma}_{it},$$

is close to $\theta := \mathbf{w}'\boldsymbol{\gamma}^\infty = \sum_{i=1}^N \sum_{t=1}^T W_{it} \Gamma_{it}^\infty$. Using the definition of $Q_{NT}(\mathbf{\Gamma}, \rho)$ and $Y_{it} = \Gamma_{it}^\infty + E_{it}$, for $D_{it} = 1$, we find that

$$L_{NT}(\theta, \rho) \leq Q_{NT}(\mathbf{\Gamma}^\infty, \rho) = \frac{1}{2} \sum_{i=1}^N \sum_{t=1}^T D_{it} E_{it}^2 + \rho \|\mathbf{\Gamma}^\infty\|_1. \tag{A.9}$$

If for a given value of $\theta = \mathbf{w}'\boldsymbol{\gamma}$ we have that the matrix $\mathbf{M}(\theta)$ with elements $M_{it}(\theta) := D_{it} E_{it} - \frac{\mathbf{w}'(\boldsymbol{\gamma} - \boldsymbol{\gamma}^\infty)}{\mathbf{w}'\text{diag}(\mathbf{p})^{-1}\mathbf{w}} \frac{(D_{it} - P_{it})W_{it}}{P_{it}}$ satisfies $\|\mathbf{M}(\theta)\|_\infty \leq \rho$, then by the definition of $\|\cdot\|_1$ in (A.8) we have $\rho \|\mathbf{\Gamma}\|_1 \leq$

$\text{Tr}(\mathbf{\Gamma}'\mathbf{M}(\theta)) = \sum_{i=1}^N \sum_{t=1}^T M_{it}(\theta)\Gamma_{it}$. Using this and $Y_{it} = \Gamma_{it}^\infty + E_{it}$, for $D_{it} = 1$, we find that

$$\begin{aligned} Q_{NT}(\mathbf{\Gamma}, \rho) &= \frac{1}{2} \sum_{i=1}^N \sum_{t=1}^T D_{it} (Y_{it} - \Gamma_{it})^2 + \rho \|\mathbf{\Gamma}\|_1 \\ &\geq \frac{1}{2} \sum_{i=1}^N \sum_{t=1}^T D_{it} [(\Gamma_{it}^\infty - \Gamma_{it}) + E_{it}]^2 + \sum_{i=1}^N \sum_{t=1}^T \left\{ D_{it} E_{it} - \frac{[(\boldsymbol{\gamma} - \boldsymbol{\gamma}^\infty)' \mathbf{w}]}{\mathbf{w}' \text{diag}(\mathbf{p})^{-1} \mathbf{w}} \frac{(D_{it} - P_{it})W_{it}}{P_{it}} \right\} \Gamma_{it} \\ &= \underbrace{\frac{1}{2} \sum_{i=1}^N \sum_{t=1}^T D_{it} (\Gamma_{it} - \Gamma_{it}^\infty)^2 - \frac{[(\boldsymbol{\gamma} - \boldsymbol{\gamma}^\infty)' \mathbf{w}]}{\mathbf{w}' \text{diag}(\mathbf{p})^{-1} \mathbf{w}} \sum_{i=1}^N \sum_{t=1}^T \frac{(D_{it} - P_{it})W_{it}}{P_{it}} (\Gamma_{it} - \Gamma_{it}^\infty)}_{=: Q_{NT}^{(\text{low},1)}(\mathbf{\Gamma})} \\ &\quad + \underbrace{\sum_{i=1}^N \sum_{t=1}^T M_{it}(\theta) \Gamma_{it}^\infty + \frac{1}{2} \sum_{i=1}^N \sum_{t=1}^T D_{it} E_{it}^2}_{=: Q_{NT}^{(\text{low},2)}} \end{aligned}$$

where in the last step we added and subtracted $\sum_{i=1}^N \sum_{t=1}^T M_{it}(\theta) \Gamma_{it}^\infty$, and we multiplied out $[(\Gamma_{it}^\infty - \Gamma_{it}) + E_{it}]^2$, which leads to some simplifications. Notice that $D_{it} E_{it} = E_{it}$ by construction of E_{it} , so that some occurrences of D_{it} above could be dropped, but we find it clearer to keep track of D_{it} explicitly here.

Next, we define the $NT \times NT$ idempotent matrices $\mathbf{P} = \frac{\text{diag}(\mathbf{p})^{-1} \mathbf{w} \mathbf{w}'}{\mathbf{w}' \text{diag}(\mathbf{p})^{-1} \mathbf{w}}$ and $\mathbf{R} = \mathbf{I}_{NT} - \mathbf{P}$. We then have

$$\begin{aligned} &Q_{NT}^{(\text{low},1)}(\mathbf{\Gamma}) \\ &= \frac{1}{2} (\boldsymbol{\gamma} - \boldsymbol{\gamma}^\infty)' \text{diag}(\mathbf{d}) (\boldsymbol{\gamma} - \boldsymbol{\gamma}^\infty) - \frac{[(\boldsymbol{\gamma} - \boldsymbol{\gamma}^\infty)' \mathbf{w}]}{\mathbf{w}' \text{diag}(\mathbf{p})^{-1} \mathbf{w}} \left[\mathbf{w}' \text{diag}(\mathbf{p})^{-1} \text{diag}(\mathbf{d} - \mathbf{p}) (\boldsymbol{\gamma} - \boldsymbol{\gamma}^\infty) \right] \\ &= \frac{1}{2} (\boldsymbol{\gamma} - \boldsymbol{\gamma}^\infty)' (\mathbf{P}' + \mathbf{R}') \text{diag}(\mathbf{d}) (\mathbf{P} + \mathbf{R}) (\boldsymbol{\gamma} - \boldsymbol{\gamma}^\infty) \\ &\quad - (\boldsymbol{\gamma} - \boldsymbol{\gamma}^\infty)' \mathbf{P}' \text{diag}(\mathbf{d} - \mathbf{p}) (\mathbf{P} + \mathbf{R}) (\boldsymbol{\gamma} - \boldsymbol{\gamma}^\infty) \\ &= \frac{1}{2} (\boldsymbol{\gamma} - \boldsymbol{\gamma}^\infty)' \mathbf{R}' \text{diag}(\mathbf{d}) \mathbf{R} (\boldsymbol{\gamma} - \boldsymbol{\gamma}^\infty) + \frac{1}{2} (\boldsymbol{\gamma} - \boldsymbol{\gamma}^\infty)' \mathbf{P}' \text{diag}(2\mathbf{p} - \mathbf{d}) \mathbf{P} (\boldsymbol{\gamma} - \boldsymbol{\gamma}^\infty), \\ &= \frac{1}{2} (\boldsymbol{\gamma} - \boldsymbol{\gamma}^\infty)' \mathbf{R}' \text{diag}(\mathbf{d}) \mathbf{R} (\boldsymbol{\gamma} - \boldsymbol{\gamma}^\infty) \\ &\quad + \frac{1}{2} (\boldsymbol{\gamma} - \boldsymbol{\gamma}^\infty)' \mathbf{P}' \text{diag}(\mathbf{p} - \mathbf{d}) \mathbf{P} (\boldsymbol{\gamma} - \boldsymbol{\gamma}^\infty) + \frac{1}{2} \frac{[(\boldsymbol{\gamma} - \boldsymbol{\gamma}^\infty)' \mathbf{w}]^2}{\mathbf{w}' \text{diag}(\mathbf{p})^{-1} \mathbf{w}}, \end{aligned}$$

where all the 'mixed terms' (that involve both \mathbf{P} and \mathbf{R}) cancel because we have $\mathbf{P}' \text{diag}(\mathbf{p}) \mathbf{R} = 0$, and in the last step we used that $\mathbf{P}' \text{diag}(\mathbf{p}) \mathbf{P} = \frac{\mathbf{w} \mathbf{w}'}{\mathbf{w}' \text{diag}(\mathbf{p})^{-1} \mathbf{w}}$. We have

$$\min_{\{\mathbf{\Gamma} \in \mathbb{R}^{N \times T} : \theta = \mathbf{w}' \boldsymbol{\gamma}\}} (\boldsymbol{\gamma} - \boldsymbol{\gamma}^\infty)' \mathbf{R}' \text{diag}(\mathbf{d}) \mathbf{R} (\boldsymbol{\gamma} - \boldsymbol{\gamma}^\infty) = 0,$$

because $\boldsymbol{\gamma}^* = \mathbf{R}\boldsymbol{\gamma}^\infty + \theta \frac{\text{diag}(\mathbf{p})^{-1}\mathbf{w}}{\mathbf{w}'\text{diag}(\mathbf{p})^{-1}\mathbf{w}}$ is a possible choice in the minimization problem, which satisfies $\mathbf{w}'\boldsymbol{\gamma}^* = \theta$ and $\mathbf{R}(\boldsymbol{\gamma}^* - \boldsymbol{\gamma}^\infty) = \mathbf{0}$. We therefore have

$$\begin{aligned} & \min_{\{\boldsymbol{\Gamma} \in \mathbb{R}^{N \times T} : \theta = \mathbf{w}'\boldsymbol{\gamma}\}} Q_{NT}^{(\text{low},1)}(\boldsymbol{\Gamma}) \\ &= \frac{1}{2} (\theta - \theta_0)^2 \left(\frac{1}{\mathbf{w}'\text{diag}(\mathbf{p})^{-1}\mathbf{w}} + \frac{\mathbf{w}'\text{diag}(\mathbf{p})^{-1}\text{diag}(\mathbf{p} - \mathbf{d})\text{diag}(\mathbf{p})^{-1}\mathbf{w}}{(\mathbf{w}'\text{diag}(\mathbf{p})^{-1}\mathbf{w})^2} \right) \\ &= \frac{1}{2} (\theta - \theta_0)^2 \left(\frac{1}{\sum_{i=1}^N \sum_{t=1}^T W_{it}^2 P_{it}^{-1}} + \frac{\sum_{i=1}^N \sum_{t=1}^T W_{it}^2 P_{it}^{-2} (P_{it} - D_{it})}{(\sum_{i=1}^N \sum_{t=1}^T W_{it}^2 P_{it}^{-1})^2} \right) \\ &= \frac{NT}{2} c_1 (v - v_0)^2, \end{aligned}$$

with c_1 as defined in the statement of the proposition, and $v - v_0 = (NT)^{-1} (\theta - \theta_0)$.

Thus, if $M_{it}(\theta) = D_{it} E_{it} - (v - v_0) V_{it}$ satisfies $\|\mathbf{M}(\theta)\|_\infty \leq \rho$, then we have

$$\begin{aligned} L_{NT}(\theta, \rho) &\geq \min_{\{\boldsymbol{\Gamma} \in \mathbb{R}^{N \times T} : \theta = \mathbf{w}'\boldsymbol{\gamma}\}} Q_{NT}^{(\text{low},1)}(\boldsymbol{\Gamma}) + Q_{NT}^{(\text{low},2)} \\ &= \frac{NT}{2} c_1 (v - v_0)^2 + \sum_{i=1}^N \sum_{t=1}^T M_{it}(\theta) \Gamma_{it}^\infty + \frac{1}{2} \sum_{i=1}^N \sum_{t=1}^T D_{it} E_{it}^2, \end{aligned}$$

and combining this with (A.9) gives

$$\begin{aligned} \frac{L_{NT}(\theta, \rho) - L_{NT}(\theta_0, \rho)}{NT} &\geq \frac{c_1}{2} (v - v_0)^2 + \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T M_{it}(\theta) \Gamma_{it}^\infty - \frac{\rho}{NT} \|\boldsymbol{\Gamma}^\infty\|_1 \\ &= \frac{c_1}{2} (v - v_0)^2 + \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T D_{it} E_{it} \Gamma_{it}^\infty \\ &\quad - (v - v_0) \frac{1}{NT} \sum_{i=1}^N \sum_{t=1}^T V_{it} \Gamma_{it}^\infty - \frac{\rho}{NT} \|\boldsymbol{\Gamma}^\infty\|_1. \end{aligned}$$

Using the assumption $c_1 > 0$ and definitions of c_2 and c_3 in the proposition, this inequality can equivalently be written as

$$\begin{aligned} \frac{2 [L_{NT}(NTv, \rho) - L_{NT}(NTv_0, \rho)]}{c_1 NT} &\geq (v - v_0)^2 - \frac{2c_2}{c_1} (v - v_0) + \left(\frac{c_2}{c_1}\right)^2 - c_3 \\ &= \left(v - v_0 - \frac{c_2}{c_1}\right)^2 - c_3. \end{aligned} \tag{A.10}$$

Notice that $c_3 > 0$ because our assumptions guarantee that $\|\mathbf{E}\|_\infty < \rho$ and therefore $\rho \|\boldsymbol{\Gamma}^\infty\|_1 \geq \sum_{i=1}^N \sum_{t=1}^T E_{it} \Gamma_{it}^\infty$, according to (A.8).

The inequality in (A.10) was derived under the assumption that $\|\mathbf{M}(NTv)\|_\infty \leq \rho$. Define $v_\pm^*(\varepsilon) \in \mathbb{R}$ and $v_\pm^*(\varepsilon) \in \mathbb{R}$ by

$$v_\pm^*(\varepsilon) := v_0 \pm (c_4 + \varepsilon), \quad \text{for } 0 < \varepsilon \leq \frac{\rho - \|\mathbf{E}\|_\infty - c_4 \|\mathbf{V}\|_\infty}{\|\mathbf{V}\|_\infty}.$$

Our assumption $\|\mathbf{E}\|_\infty + c_4 \|\mathbf{V}\|_\infty < \rho$ guarantees that such an $\varepsilon > 0$ exists. Using the triangle inequality we find that

$$\|\mathbf{M}(NTv_\pm^*(\varepsilon))\|_\infty = \|\mathbf{E} - (v_\pm^*(\varepsilon) - v_0)\mathbf{V}\|_\infty \leq \|\mathbf{E}\|_\infty + |v_\pm^*(\varepsilon) - v_0| \|\mathbf{V}\|_\infty \leq \rho,$$

where the final inequality follows from the definition of $v_{\pm}^*(\varepsilon)$. The conditions for (A.10) is therefore satisfied by $v = v_{\pm}^*(\varepsilon)$; that is, we have

$$\begin{aligned} \frac{2 [L_{NT}(NT v_{\pm}^*(\varepsilon), \rho) - L_{NT}(NT v_0, \rho)]}{c_1 NT} &\geq \left(v_{\pm}^*(\varepsilon) - v_0 - \frac{c_2}{c_1} \right)^2 - c_3 \\ &= \left(c_4 + \varepsilon \mp \frac{c_2}{c_1} \right)^2 - c_3 \\ &= \left(\sqrt{c_3} + \varepsilon + \frac{|c_2| \mp c_2}{c_1} \right)^2 - c_3 \\ &\geq (\sqrt{c_3} + \varepsilon)^2 - c_3 \\ &> 0, \end{aligned}$$

where we used the definition $c_4 = \sqrt{c_3} + \frac{|c_2|}{c_1}$.

$L_{NT}(NT v, \rho)$ is a convex function of $v = \theta/NT$, because it was obtained via profiling of the convex function $Q_{NT}(\Gamma, \rho)$. The value v_0 lies in the interval $[v_+^*(\varepsilon), v_-^*(\varepsilon)]$, and we have shown that $L_{NT}(NT v_0, \rho) < L_{NT}(NT v_{\pm}^*(\varepsilon), \rho)$. It must therefore be the case that the optimal $\widehat{v} = NT\widehat{\theta}$ that minimizes $L_{NT}(NT v, \rho)$ also lies in the interval $[v_+^*(\varepsilon), v_-^*(\varepsilon)]$ —otherwise we obtain a contradiction to the convexity of $L_{NT}(NT v, \rho)$. Thus, we have shown that

$$|\widehat{v} - v_0| \leq c_4 + \varepsilon,$$

and because we can choose $\varepsilon > 0$ arbitrarily small it must be the case that

$$|\widehat{v} - v_0| \leq c_4,$$

which is what we wanted to show. □