Disaggregated impulse response functions via the classifier-Lasso*

Miguel Bandeira
Insper and CFM

Igor Martins
Örebro University

Abstract

The commonplace approach to the estimation of disaggregated impulse responses functions involves ex-ante grouping of individuals and subsequently estimating the associated group-specific responses. This paper shows that impulse response estimates based on this approach are subject a misclassification bias that arises whenever researchers groups together individuals that do not react in the same way to shocks. This paper proposes a methodology to estimate disaggregated impulse response functions using the classifier-Lasso which asymptotically eliminates the misclassification bias. The methodology is used to estimate the dynamic responses of firm-level debt to an aggregate investment specific technology shock.

JEL Classification: C22, C23, C38, E22.

*Contact details: Bandeira (mbandeira@insper.edu.br) and Martins (igor.ferreira-batista-martins@oru.se. For insightful comments and discussions, the authors would like to thank Andrea Alati, Regis Barnichon, Adrien Bussy, Thomas Drechsel, Mike Elsby, Friedrich Geiecke, Diogo Guillén, Daniel Lewis, Ben Moll, Lukasz Rachel, Ricardo Reis and seminar participants at the 2021 Luso-Brazilian Macroeconomics Meeting (LubraMacro). Any remaining errors are our own. This version: 9th July, 2025.

1. Introduction

Ever since the seminal contributions of Frisch (1933) and Slutzky (1937), macroeconomists have looked at random shocks as the primary source of business cycle fluctuations (Ramey, 2016). For most of the twentieth century, the bulk of the work in empirical macroeconomics was devoted to identification of macroeconomic shocks and estimation of impulse response functions (IRFs) which summarised the expected dynamic responses of aggregate variables to these shocks. More recently, with the increasing availability of micro data and growing popularity of heterogeneous agent models, the focus of the literature has gradually shifted towards understanding how different firms and households respond to these aggregate macroeconomic shocks. The estimation of *disaggregated IRFs* is important not only to quantify what drives the responses of aggregate variables but also as a means to empirically test the predictions from alternative transmission theories.

This paper studies the estimation of disaggregated IRFs in a setting with *latent group heterogeneity*. Our setting is characterised by three main features: (*i*) each individual belongs to a group within a broadly heterogeneous population, (*ii*) the individual IRFs are the same *within* a group but differ *between* groups and (*iii*) the researcher has disaggregated data on the outcome interest and a strictly exogenous macroeconomic shock, but does *not* know either how many groups are in the population or which group each individual belongs to. This setting is illustrated in figure 1. The estimated individual-specific IRFs (grey lines) are obtained from a dataset in which half of the individuals belong to a group for which the true IRF is given by the green-solid line whilst the other half the true IRF is given by the red-solid line. The problem considered in this paper is that of a researcher that observes a dataset that yields "cloud" of grey-lines and, based on that dataset, would like to estimate: (*i*) the number of latent groups in the population, (*ii*) the associated group-specific IRFs and (*iii*) which individuals belong to which group.

In the existing literature, the common approach to the estimation of disaggregated IRFs involves first grouping individuals in the sample according to some external classification or observable explanatory variables and subsequently estimating the group-specific IRFs by pooling together individuals that are assumed to belong to the same group. For example, when studying the responses of consumption expenditures to monetary policy changes households can be grouped according to their housing tenure status (Cloyne, Ferreira and Surico, 2019), according to their age (Wong, 2019) or according to their relative position in the wealth distribution (Coibion, Gorodnichenko, Kueng and Silvia, 2017). This paper shows that, in the presence of latent group heterogeneity, this ex-ante grouping approach can lead to misleading conclusions. More precisely, it is shown that there is a bias-variance tradeoff between the IRF estimates obtained by ex-ante grouping and the IRFs estimated for each individual separately. The IRF estimates based on the ex-ante grouping of individuals are more precise but are subject to a form of bias, which is here labeled *misclassification bias*, that arises whenever the grouping of individuals imposed by the researcher groups together individuals that do not react in the same way to aggregate shocks.

Motivated by this theoretical result, this paper introduces an alternative methodology to estimate disaggregated IRFs in the presence of latent group heterogeneity. The methodology builds from penalized estimation techniques, in particular the classifier-Lasso (C-Lasso) from Su, Shi and Phillips (2016), in order to simultaneously estimate the unknown group-specific IRFs and classify individuals to groups whereas the number of latent groups is subsequently estimated via a BIC-type information criterion. The theoretical results in Su, Shi and Phillips (2014, 2016) imply that in our setting the group-specific IRFs estimator based on the C-Lasso have the same asymptotic properties as the *group-oracle* IRFs, that is, the IRF estimates that would result from an ex-ante grouping of individuals that exactly matches the true unknown grouping of individuals. Most importantly, and in sharp contrast with ex-ante grouping approach, the C-Lasso IRF estimator achieve this property in a completely data-driven way that does not require the researcher to take a stance on either the number of latent groups or the individual group membership.

To illustrate the finite sample performance of the C-Lasso based classification and estimation procedure, this paper uses a Monte Carlo experiment in which artificial datasets are generated from the same DGP used to generate the IRFs in figure 1. In each Monte Carlo sample the C-Lasso framework is used to estimate the IRFs both by estimating the whole moving average representation directly and by local projections (Jordà, 2005). The results from this Monte Carlo experiment illustrate the good performance of the C-Lasso in terms of determination of the number of latent groups, classification of individuals into different groups and estimation of group-specific IRFs in samples of similar size than those typically used in the existing literature estimating disaggregated IRFs.

As an empirical application, the proposed C-Lasso framework is used to revisit the dynamic responses of firm-level debt to an aggregate investment specific technology (IST) shock from Drechsel (2023). A theoretical prediction from the model in Drechsel (2023) is that the debt of firms that tend to borrow against collateral should decrease following a positive IST shock whereas the debt of firms that tend to borrow against their future earnings should increase. When applied to a subset of the Drechsel (2023) dataset, the C-Lasso framework identifies two latent groups, one for which the response of firm-level debt to an IST shock is positive and other for which debt reacts negatively. The group of firms that increase their debt following an IST shock is composed of firms that are relatively smaller, have a higher share of intangible assets, tend to be earnings-based borrowers and do not belong to the consumer staples or utilities sectors. Altogether these findings are in line with the theoretical predictions from Drechsel (2023), but also suggest that, on top of whether a firm tends to borrow against earnings or collateral, the specific *sector* that the firm operates also plays a role in determining whether it will increase or decrease its debt in response to aggregate IST shocks.

Relation to the literature. This paper relates and contributes to two strands of literature. First and foremost, it relates to the empirical macroeconomics literature that focuses on the estimation of impulse response functions and, in particular, to several empirical applications that estimate heterogeneous impulse responses functions to aggregate shocks (see section 2.2 for a review of some of these applications). This paper contributes to this literature in two fundamental ways. First, by formally showing that there is a bias-variance tradeoff between the estimator based on ex-ante individual classification and the estimator based on individual-specific impulse responses. Second, by introducing an alternative methodology that produces estimates that achieve a smaller mean squared error without the need to ex-ante take a stance on individual group membership.

Second, it relates to an extensive literature on variable-coefficient models in panel data (see, for instance, Hsiao, 2014, chapter 6) and, in particular, to panel structure models where individuals are assumed to belong to a number of homogeneous groups within a broadly heterogeneous population and the regression parameters are the same *within* each group but differ *across* groups. Different approaches have been proposed to determine an unknown group structure in modeling unobserved slope heterogeneity in panels, including finite mixture models (e.g. Sun, 2005; Kasahara and Shimotsu, 2009; Browning and Carro, 2014), variants of the K-means algorithm (e.g. Lin and Ng, 2012; Sarafidis and Weber, 2015; Bonhomme and Manresa, 2015) and penalized estimation techniques (e.g. Su, Shi and Phillips, 2014, 2016; Wang, Phillips and Su, 2018).

Most closely related to the present paper are two contemporaneous works that also build on this literature to estimate effects of shocks on different individuals. First, Lewis, Melcangi and Pilossoph (2022) use a Gaussian mixture linear regression to estimate the distribution of marginal propensities to consume using the 2008 tax rebate in the United States. Second, **?** introduces group local projections which build from the K-means algorithm to group heterogenous IRFs within a local projection-IV framework. This paper is the first to use penalized estimation techniques, and in particular the C-Lasso, to estimate disaggregated impulse response functions. In addition to the desirable asymptotic properties of the proposed estimator that follow from the theoretical results in Su, Shi and Phillips (2016), a Monte Carlo experiment is provided to illustrate the good finite sample performance of the proposed estimator with sample sizes commonly used by the existing literature estimating disaggregated IRFs.

Structure of the paper. Section 2 introduces the data generating process and discusses some of the empirical applications it can accommodate. Section 3 discusses the statistical properties of the common approach used to estimate heterogeneous impulse responses in the literature. Section 4 introduces a C-Lasso based methodology to estimate heterogeneous impulse responses in the presence of latent group heterogeneity and discusses its asymptotic properties. Section 5 uses a Monte Carlo experiment to illustrate the finite sample properties of the proposed methodology. Section 6 applies this methodology to revisit the Drechsel (2023) impulse response estimates of firm level debt to an aggregate IST shock. Section 7 concludes and discusses avenues for future research.

Notation. In all that follows, bold letters are used to denote vectors or matrices and non-bold fonts denote scalars. For a given matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$, $\mathbf{A}'$ denotes its transpose, $\mathbf{a}_{i,*}$ denotes its i-th row, $\mathbf{a}_{*,j}$ denotes its j-th column and $a_{i,j}$ denotes its (i, j)-th element. Moreover, $\mathbf{1}_{m \times n}$ and $\mathbf{0}_{m \times n}$ denote $m \times n$ matrices of ones and zeros, $\mathbb{1}\{\cdot\}$ denotes the indicator function, $\|\cdot\|$ denotes the Frobenius norm, $\otimes$ denotes the Kronecker product and $\oplus$ the direct sum of matrices. For any two real numbers $a < b$, denote by $\mathbb{Z}_{[a,b]}$ the set of all integers in $[a, b]$.

2. Impulse response functions under latent group heterogeneity

This section starts by introducing the process that is assumed to generate the panel data set observed by the researcher. Then, it discusses how it can accommodate some specifications that have been used in the literature to estimate heterogenous impulse response functions.

2.1. The data generating process

In a nutshell, the assumed data generating process can be characterised as a distributed lag model in which the coefficients are allowed to vary across different groups of individuals. The distributed lag model assumption can be formalised as follows,

***Assumption* 1** *The researcher observes a panel data set* $\{(y_{i,t}, \mathbf{x}_{i,t})\}$ *for* $i = 1, \ldots, N$ *and* $t = -H + 1, \ldots, -1, 0, 1, \ldots, T$ *for which the data generating process can be represented as,*

$$y_{i,t} = \mathbf{x}'_{i,t} \boldsymbol{\beta}_i + \varepsilon_{i,t} \tag{1}$$

where $\mathbf{x}'_{i,t} \equiv [x_{i,t}, x_{i,t-1}, \ldots, x_{i,t-H}]$ *and* $\boldsymbol{\beta}_i = [\beta_{i,0}, \ldots, \beta_{i,H}]$. *Moreover, let* $\mathbf{X}_i = [\mathbf{x}_{i,1}, \ldots, \mathbf{x}_{i,T}]'$ *and* $\mathbf{X} = \{\mathbf{X}_i\}_{i=1}^{n}$ *and assume the following three conditions hold:* (a) *rank* $(\mathbf{X}'_i\mathbf{X}_i) = H + 1$, $\forall i$; (b) $\mathbb{E}(\varepsilon_{i,t} \mid \mathbf{X}) = 0$, $\forall i, t$ *and* (c) $\mathbb{C}\text{ov}(\varepsilon_{i,s}, \varepsilon_{j,t} \mid \mathbf{X}) = \mathbb{1}\{i = j\}\mathbb{1}\{s = t\}\,\sigma^2$.

Assumption 1 is a panel version of a distributed lag model.[2] For the macroeconomic applications that are the focus of the present paper, the independent variables typically consist of a macroeconomic shock and its lags, for instance, a monetary policy shock or an aggregate productivity shock, in which case (1) is a finite

[2] See, for instance, Greene (2003, chapter 19) or Baltagi (2008, chapter 6).

moving average representation and the vector $\boldsymbol{\beta}_i$ is the impulse response function of the i-th individual in the panel to that shock.[3] In that context, it is assumed that the researcher observes a panel containing the dependent variable of interest for different individuals over time and a time-series of the macroeconomic shock of interest.[4] The error term in (1) captures any other individual-specific factors that affect the dependent variable of interest and it is assumed those factors are mean independent of the macroeconomic shock and conditionally homoskedastic.[5] In order to ensure the estimators considered in this paper are well-defined, it is further assumed that the macroeconomic shock is not perfectly collinear with any of its lags.

In addition to assumption 1, the individual impulse response functions to the macroeconomic shock are assumed to follow a group pattern of the form,

***Assumption* 2** *The individual specific coefficients in* (1) *follow a group pattern of the form,*

$$\boldsymbol{\beta}_i = \sum_{k=1}^{K_0} \boldsymbol{\alpha}_k \mathbb{1}\left\{i \in G_k\right\} \tag{2}$$

where $\boldsymbol{\alpha}_j \neq \boldsymbol{\alpha}_k$ *for any* $j \neq k$, $\cup_{k=1}^{K_0} G_k = \{1, 2, \ldots, N\}$ *and* $G_j \cap G_k = \varnothing$ *for any* $j \neq k$.

[3] There can be different data generating processes that can be represented by a moving average representation. For instance, it can arise from the inversion of a panel vector autoregression or an autoregressive distributed lag model. If that is the case, the implicit assumption in (1) is that the coefficients of the autoregressive component are such that that the process admits a linear moving average representation in which the moving average coefficients decay sufficiently fast to ensure that they can be truncated at a finite horizon H.

[4] Over the last three decades, the empirical macroeconomics literature has come up with a wide range of methods to identify macroeconomic shocks, including several identification schemes for structural VARs, narrative identification (e.g., Romer and Romer, 2004) and high-frequency identification (e.g., Gertler and Karadi, 2015). Different methods of identifying macroeconomic shocks are surveyed in Ramey (2016).

[5] For expositional purposes it is useful to focus on the parsimonious moving average representation in (1). Notice, however, that all the IRF estimators presented in this paper are either based on OLS or Penalized least squares the discussion could be extended to include additional independent variables in (1) including a constant or individual fixed-effects. In those cases, one would simply need to use residualized versions of the dependent and independent variables.

Assumption 2 imposes the same form of coefficient heterogeneity that is assumed in Su, Shi and Phillips (2016) and, in the present context, can be justified by the idea that there is some form of *group sparsity* in the way different individuals react to macroeconomic shocks. Put differently, assumption 2 represents a middle-ground between two extreme scenarios. One in which every single individual reacts in a different way to a macroeconomic shock and other where all individuals react in the same way to that shock. Instead, according to (2), individuals can belong to one among K_0 groups and the impulse responses differ between groups but are common across individuals within the same group. This assumption maps naturally into macroeconomic models featuring *ex-ante heterogeneity*, for example, models with optimizing and hand-to-mouth consumers (e.g. Galí, López-Salido and Vallés, 2007; Bilbiie, 2008) or models in which the degree of price rigidity differs across sectors (e.g. Carvalho, 2006). However, because it imposes that impulse responses of each group do not vary over time and that individuals cannot switch between groups over time, our setting does not map into models where individual impulse responses vary over time such as models featuring aggregate shocks and *ex-post heterogeneity* due to uninsurable idiosyncratic shocks.

Finally, it is important to notice that from the researcher's perspective both the true number of groups and which individuals belong to which group are *unknown* and, therefore, the problem of estimating individual impulse responses and understanding what drives their heterogeneity is equivalent to estimating the number of groups (K_0), the individual group membership ($\{G_1, \ldots, G_{K_0}\}$) and the group specific impulse responses ($\boldsymbol{\alpha}_1, \ldots, \boldsymbol{\alpha}_{K_0}$).

2.2. Empirical applications

Before turning to the estimation of individual specific impulse responses in the presence of latent group heterogeneity, it is useful to review some heterogeneity analysis conducted in the existing literature both to illustrate some of the settings where the methods developed in this paper could be applied to and to understand what is the “common approach” in the literature to estimate the responses of different individuals to a common aggregate shock.

Heterogeneous impulse responses to a monetary policy shock. There is an extensive list of heterogeneity analysis that have been conducted to investigate to what extent monetary policy surprises have heterogeneous effects across different individuals, firms or regions. For instance, Coibion, Gorodnichenko, Kueng and Silvia (2017) investigate the effects of monetary policy shocks on consumption of individuals depending on their relative position in the wealth distribution, Wong (2019) investigates whether monetary policy affects differently consumption expenditures of households depending on their age whilst Cloyne, Ferreira and Surico (2019) analyze the responses of consumption expenditures depending on whether the household is a home owner, a renter or a mortgagor. Moreover, some papers have investigated the responses of inflation based on group specific inflation baskets, for instance, Cravino, Lan and Levchenko (2020) investigate the effects of monetary policy on the inflation experienced by individuals in different percentiles of the income distribution whereas Clayton, Jaravel and Schaab (2018) find that monetary policy stabilizes sectors that matter relatively more for college-educated households. Bernanke and Gertler (1995) investigate the effects of monetary policy on different components of final demand, Carlino and Defina (1999) investigate the effects of monetary policy on the state-level economic activity across US states.

Numerous papers have also looked at the effects of a monetary policy shock across different types of firms. For instance, Gertler and Gilchrist (1994) find that small firms account for a significantly disproportionate share of the manufacturing declines that follows tightening of monetary policy, Kashyap, Lamont and Stein (1994) find that during the 1981-82 recession bank-dependent liquidity constrained firms cut their inventories by significantly more than their nonbank-dependent counterparts. More recently, Ottonello and Winberry (2020) find that the investment of firms with low default risk is the most responsive to monetary shocks, Jeenas (2019) finds that are the firms with fewer liquid assets that tend to reduce investment relative to others and Cloyne, Ferreira, Froemel and Surico (2023) find that younger firms paying no-dividends adjust both their capital expenditure and borrowing significantly more than older firms paying dividends in response to a monetary policy shock.

Other shocks. In the same spirit, heterogeneity analysis have been conducted to understand the reactions of different groups of individuals or firms to other aggregate shocks. In particular, Drechsel (2023) investigates the effects of an aggregate investment specific technology shock on firm level debt depending on whether firms tend to borrow against their collateral or future earnings. This application will be revisited in section 6.

3. The common approach to estimation of heterogeneous impulse responses

Even though existing heterogeneity analyses focus different dimensions, methodologically they mostly follow an *ex-ante classification approach*, that is, they first group individuals according to some external classification or observable explanatory variables and then estimate and compare the resulting group specific impulse responses. This section analyses the properties of the estimator based on the ex-ante classification approach and shows that in the presence of latent group heterogeneity there is, in general, a *bias-variance tradeoff* between this estimator and estimating individual-specific impulse responses.

3.1. Ex-ante classification and individual-specific impulse responses

To introduce the estimator based on the ex-ante classification approach, let a *grouping scheme* be denoted by $\tilde{\mathcal{G}}^K$ which stands for any collection of K non-empty sets satisfying $\cup_{k=1}^{K} \tilde{G}_k = \{1, 2, \ldots, N\}$ and $\tilde{G}_i \cap \tilde{G}_j = \varnothing$ for any $i \neq j$. In practice, the choice of individual group membership might be a function of other observable variables (*e.g.* the income distribution decile that the individual belongs, the household house-tenure status or whether a firm is a flow or a collateral borrower). Most importantly, $\tilde{\mathcal{G}}^K$ is a researcher's choice and both the number of groups chosen and the individual classification can differ from the true number of groups and group membership in assumption 2. Given a grouping scheme $\tilde{\mathcal{G}}^K$, the impulse response estimator based on the ex-ante classification approach is defined by,

$$\widetilde{\boldsymbol{\beta}}_i(\tilde{\mathcal{G}}^K) = \sum_{k=1}^{K} \widetilde{\boldsymbol{\alpha}}_k \mathbb{1}\left\{i \in \tilde{G}_k\right\} \tag{3}$$

where,

$$\left(\widetilde{\boldsymbol{\alpha}}_1(\tilde{\mathcal{G}}^K), \ldots, \widetilde{\boldsymbol{\alpha}}_K(\tilde{\mathcal{G}}^K)\right) = \underset{\mathbf{a}_1, \ldots, \mathbf{a}_K}{\arg\min} \frac{1}{NT} \sum_{i=1}^{N} \sum_{t=1}^{T} \left(y_{i,t} - \mathbf{x}_{i,t}' \sum_{k=1}^{K} \mathbf{a}_k \mathbb{1}\{i \in \tilde{G}_k\}\right)^2 \tag{4}$$

The group estimates obtained from (4) are nothing more than ordinary least squares estimates obtained from pooling all the individuals in the panel and interacting the shock with a dummy variable for group membership. Once the group estimates are obtained, (3) uses the grouping scheme to assign impulse response estimates to each individual in the panel.

As a benchmark, it will be useful to consider the estimator of impulse responses that would be obtained if the researcher did not take a stance on the grouping scheme and instead allowed for complete heterogeneity in the individual responses to the shock. That estimator is defined by,

$$\left(\widehat{\beta}_1, \ldots, \widehat{\beta}_N\right) = \underset{\mathbf{b}_1,\ldots,\mathbf{b}_N}{\arg\min} \ \frac{1}{NT} \sum_{i=1}^{N} \sum_{t=1}^{T} \left(y_{i,t} - \mathbf{x}_{i,t}' \mathbf{b}_i\right)^2 \tag{5}$$

where the individual-specific impulse response estimates obtained from (5) are the same one would obtain by estimating the moving average representation using the time-series for each individual in the panel separately.

3.2. A bias-variance tradeoff

In the presence of latent group heterogeneity the choice between the estimator based on the ex-ante classification approach and the estimator allowing for complete heterogeneity entails a bias-variance tradeoff. This tradeoff is formalised by the following proposition,

***Proposition* 1** *Suppose assumptions 1 and 2 hold. For a given $\tilde{\mathcal{G}}^K$ suppose $i \in \tilde{G}_a \cap G_b$ for some $a \in \mathbb{Z}_{[1,K]}$ and $b \in \mathbb{Z}_{[1,K^0]}$. Let $\widetilde{\beta}_i(\tilde{\mathcal{G}}^K)$ denote the estimator obtained from (3) and (4) and $\widehat{\beta}_i$ denote the estimator obtained from (5). Then,*

$$\mathbb{E}\left(\widehat{\boldsymbol{\beta}}_i\right) = \boldsymbol{\beta}_i \tag{6}$$

$$\mathbb{E}\left(\widetilde{\boldsymbol{\beta}}_i(\tilde{\mathcal{G}}^K)\right) = \varphi_{a,b}\, \boldsymbol{\beta}_i + \sum_{\substack{k=1 \\ k \neq b}}^{K^0} \varphi_{a,k}\, \boldsymbol{\alpha}_k \tag{7}$$

where $\varphi_{a,b} \equiv \mathbb{E}\left(\left(\sum_{i \in \tilde{G}_a} \mathbf{X}_i'\mathbf{X}_i\right)^{-1} \sum_{i \in \tilde{G}_a \cap G_b} \mathbf{X}_i'\mathbf{X}_i\right)$. Moreover, for any non-zero $H+1$-dimensional vector $\mathbf{r}$ it holds that,

$$\mathbf{r}'\, \mathbb{V}\text{ar}\left(\widehat{\boldsymbol{\beta}}_i \mid \mathbf{X}\right) \mathbf{r} \ \geqslant \ \mathbf{r}'\, \mathbb{V}\text{ar}\left(\widetilde{\boldsymbol{\beta}}_i(\tilde{\mathcal{G}}^K) \mid \mathbf{X}\right) \mathbf{r} \tag{8}$$

Proof. See appendix A. □

In simple terms, proposition 1 states that the researcher faces a fundamental tradeoff when deciding how to estimate individual impulse response functions in the presence of latent group heterogeneity. On the one hand, disaggregating too much could yield a set of estimated impulse responses that are largely uninformative once the variability across individuals reflects not only the latent group heterogeneity but also a large share of sampling variability. On the other hand, grouping together individuals that do not share the same responses lead to biased impulse response estimates. This tradeoff is illustrated in figure 1. The estimation of fully heterogeneous impulse responses yields the grey cloud of responses from which it is almost impossible to infer the true heterogeneity pattern which comes from the fact that half of the individuals in the sample have their true impulse responses given by the green line whereas the other half have their impulse response given by the red line. On the other hand, if individuals were incorrectly grouped and the ex-ante classification approach was adopter then patterns of heterogeneity could be mistakenly inferred from the data. For example, if all individuals where grouped together would yield the wrong conclusion that the individuals true impulse response is given by the black-dashed line and the cloud of individual impulse responses is the result of sampling variability and not the result of heterogeneity in the true impulse responses.

Misclassification bias. For a given individual i, from (7) there is only one case where $\widetilde{\beta}_i(\tilde{\mathcal{G}}^K)$ is not biased: when all the individuals assigned to the same group as i by the researcher indeed belong to the same latent group as individual i. This implies that for the ex-ante classification approach to yield unbiased estimates for *all* individuals in the sample requires that the ex-ante grouping of individuals proposed by the researcher exactly matches the the true individual grouping in assumption 2. If this is not the case, the estimator based on the ex-ante classification approach suffers from *misclassification bias*. Expression in (7) states that on average the impulse responses estimated for a given individual are equal to a weighted average between the impulse response of the group that individual belongs to and the impulse responses of other groups. For the case where $x_{i,t}$ is an aggregate shock, expression (7) becomes,

$$\mathbb{E}\left(\widetilde{\boldsymbol{\beta}}_i(\tilde{\mathcal{G}}^K)\right) = \frac{N_{\tilde{G}_a \cap G_b}}{N_{\tilde{G}_a}}\,\boldsymbol{\beta}_i + \sum_{\substack{k=1 \\ k \neq b}}^{K^0} \frac{N_{\tilde{G}_a \cap G_k}}{N_{\tilde{G}_a}}\,\boldsymbol{\alpha}_k \tag{9}$$

where $N_{\tilde{G}_a \cap G_b}$ denotes the cardinality of the set $\tilde{G}_a \cap G_b$. In this case, the weight assigned to each latent group true impulse response is given by the share of individuals from that group that where assigned by the researcher to the same group as individual i.

Scope for efficiency gains. Given the ex-ante classification approach is prone to suffer from a misclassification bias, a natural question is whether the gains in precision obtained by ex-ante grouping of individuals are sufficiently large to outweigh the risk of ending up with biased estimates. According to (8) at any horizon considered the sampling variance of the impulse responses obtained from the ex-ante classification approach have smaller (or equal) variance than their counterparts obtained from estimating fully heterogeneous impulse response estimates. For the case where $x_{i,t}$ is an aggregate shock it can be shown that,

$$\mathbb{V}\text{ar}\left(\widetilde{\boldsymbol{\beta}}_i(\tilde{\mathcal{G}}^K) \mid \mathbf{X}\right) = \frac{1}{N_{\tilde{G}_a}}\mathbb{V}\text{ar}\left(\widehat{\boldsymbol{\beta}}_i \mid \mathbf{X}\right) \tag{10}$$

where $N_{\tilde{G}_a}$ denotes the cardinality of the set $\tilde{G}_a$. In other words, by grouping together ten individuals, a relatively small number relative to the typical cross-sectional dimension in datasets used to estimate heterogeneous impulse responses, the sampling variance of impulse responses decreases by 90% which suggest that, in general, the efficiency gains from pooling individuals together can be sizable.

4. Impulse response estimation via the classifier-Lasso

Motivated by the bias-variance tradeoff in proposition 1, this section introduces an alternative way to estimate group-specific impulse responses that is designed to eliminate this tradeoff *without* the requirement that the researcher correctly specifies ex-ante the group membership. The estimation is an application of the classifier-Lasso (C-Lasso) developed in Su, Shi and Phillips (2014, 2016) to estimate group-specific impulse response functions in the presence of latent heterogeneity. The fundamental insight underlying the C-Lasso is that it builds on penalized techniques to replace ex-ante classification of individuals into groups by a data-driven way of estimating both the individual group membership and the number of latent groups. This section briefly reviews the C-Lasso and shows how it can be applied to the estimation of heterogeneous impulse responses.

4.1. Determination of individual group membership

Consider the problem of determining individual group membership taking the number of latent groups as given. First, define the following objective function (Su, Shi and Phillips, 2014, equation 2.4),

$$\mathcal{Q}^{(K)}_{NT,\lambda_1}(\mathbf{b},\mathbf{a}) = \frac{1}{NT}\sum_{i=1}^{N}\sum_{t=1}^{T}\left(y_{i,t} - \mathbf{x}'_{i,t}\mathbf{b}_i\right)^2 + \frac{\lambda_1}{N}\sum_{i=1}^{N}\prod_{k=1}^{K}\|\mathbf{b}_i - \mathbf{a}_k\| \qquad (11)$$

where λ_1 is a tunning parameter that converges to zero as $(N, T) \to \infty$. Notice that the first term on the right-hand-side of (11) is exactly the same objective function that is used to obtain impulse response estimates under complete individual heterogeneity in (5). The second-term on the right-had-side of (11) is the distinctive feature of the C-Lasso and its mixed additive-multiplicative form shrinks the individual impulse responses ($\mathbf{b}_i$) to a particular unknown group-level parameter vector ($\mathbf{a}_k$). Minimising (11) with respect to $\mathbf{b}$ and $\mathbf{a}$ produces the C-Lasso estimates $\boldsymbol{\beta}$ and $\boldsymbol{\alpha}$ which are henceforth denoted by $\widehat{\boldsymbol{\beta}}^{\text{C-Lasso}}$ and $\widehat{\boldsymbol{\alpha}}^{\text{C-Lasso}}$.

Given this set of estimates, the C-Lasso group classifier is given by: $i \in \widehat{G}_k$ if $\widehat{\boldsymbol{\beta}}_i^{\text{C-Lasso}} = \widehat{\boldsymbol{\alpha}}_k^{\text{C-Lasso}}$ for some $k \in \mathbb{Z}_{[1,K]}$, otherwise, $i \in \widehat{G}_l$ for some $l \in \mathbb{Z}_{[1,K]}$ if $||\widehat{\boldsymbol{\beta}}_i^{\text{C-Lasso}} - \widehat{\boldsymbol{\alpha}}_l^{\text{C-Lasso}}|| \leqslant \{ ||\widehat{\boldsymbol{\beta}}_i^{\text{C-Lasso}} - \widehat{\boldsymbol{\alpha}}_1^{\text{C-Lasso}}||, \ldots, ||\widehat{\boldsymbol{\beta}}_i^{\text{C-Lasso}} - \widehat{\boldsymbol{\alpha}}_K^{\text{C-Lasso}}|| \}$ and $\sum_{k=1}^{K} \mathbb{1}\{\widehat{\boldsymbol{\beta}}_i^{\text{C-Lasso}} = \widehat{\boldsymbol{\alpha}}_k^{\text{C-Lasso}}\} = 0$.[6]

4.2. Determination of the number of groups

The C-Lasso estimates from minimising (11) are obtained for a given number of groups (K). In practice, however, the true number of latent groups is unknown and has to be estimated along with the group membership. Following Su, Shi and Phillips (2014, 2016) it is assumed that the true number of groups is bounded from above by a finite integer $K_{\max}$ and the number of groups is estimated through an information criterion (IC).[7] Making the dependence on K and λ_1 explicit, the group classification implied by the C-Lasso can be written as $\widehat{G}(K, \lambda_1) = \left\{ \widehat{G}_1(K, \lambda_1), \ldots, \widehat{G}_K(K, \lambda_1) \right\}$. The information criterion used to determine the number of latent groups is given by Su, Shi and Phillips (2016, equation 2.10),

$$\text{IC}(K, \lambda_1) = \ln\left(\hat{\sigma}^2_{\widehat{G}(K,\lambda_1)}\right) + \rho_{NT}\,(H+1)\,K \tag{12}$$

where ρ_{NT} is a tuning parameter and $\hat{\sigma}^2_{\widehat{G}(K,\lambda_1)} = \frac{1}{NT} \sum_{k=1}^{K} \sum_{i \in \widehat{G}(K,\lambda_1)} \sum_{t=1}^{T} \left(y_{i,t} - \mathbf{x}'_{i,t} \widehat{\boldsymbol{\alpha}}_{\widehat{G}_k} \right)^2$ with $\widehat{\boldsymbol{\alpha}}_{\widehat{G}_k} = \left(\sum_{i \in \widehat{G}_k} \sum_{t=1}^{T} \mathbf{x}_{i,t} \mathbf{x}'_{i,t} \right)^{-1} \left(\sum_{i \in \widehat{G}_k} \sum_{t=1}^{T} \mathbf{x}_{i,t} y_{i,t} \right)$. Finally, for a given value of the tunning parameter λ_1, the number of groups is chosen such that the IC in (12) is minimized, that is, $\hat{K}(\lambda_1) = \arg\min_{1 \leqslant k \leqslant K_{max}} \text{IC}(k, \lambda_1)$.

[6] This group classifier achieves in large samples the same properties as the simpler classification rule $\widehat{G}_k = \{i \in \mathbb{Z}_{[1,N]} : \widehat{\boldsymbol{\beta}}_i^{\text{C-Lasso}} = \widehat{\boldsymbol{\alpha}}_k^{\text{C-Lasso}}\}$ for $k \in \mathbb{Z}_{[1,K]}$. Nonetheless, the classifier in the text is preferred since it ensures that all the individuals are classified into one of the K groups in finite samples (see Su, Shi and Phillips, 2016, remark 2).

[7] An alternative way of determining the number of latent groups is to use the residual-based Lagrange multiplier-type test proposed by Lu and Su (2017).

4.3. Post-Lasso impulse responses

Given the estimated group classification based on the C-Lasso this paper focuses on the *post-Lasso* estimates of the impulse responses. For a given group $\widehat{G}_k$ the post-Lasso group impulse response estimates are given by $\widehat{\boldsymbol{\alpha}}_{\widehat{G}_k}$ and the post-Lasso individual impulse responses are given by $\widehat{\boldsymbol{\beta}}_{i,\widehat{G}_k} = \sum_{k=1}^{K} \widehat{\boldsymbol{\alpha}}_{\widehat{G}_k} \mathbb{1}\left\{i \in \widehat{G}_k\right\}$.

4.4. Asymptotic properties

The asymptotic properties of the C-Lasso in the context of linear models are formally shown in Su, Shi and Phillips (2014, 2016). Under a suitable set of assumptions, the authors show that: (*i*) the classifier proposed in section 4.1 is *uniformly consistent* which, in simple terms, means that the proposed C-Lasso group classifier classifies each individual to the correct group with probability approaching 1 as $(N, T) \to \infty$ (see Su, Shi and Phillips, 2016, theorem 2.2); (*ii*) the selector criterion for K proposed in section 4.2 is such that $\mathbb{P}\left(\hat{K}(\lambda_1) = K_0\right) \to 1$ as $(N, T) \to \infty$ (see Su, Shi and Phillips, 2016, theorem 2.6) and (*iii*) the post-Lasso estimator of $\boldsymbol{\alpha}_k$ defined in section 4.3 enjoy the *asymptotic oracle property*, in particular, as $(N, T) \to \infty$ it achieves the same limiting distribution as the *oracle estimator* which is the group impulse response estimator one would obtain if the true group membership was known (see Su, Shi and Phillips, 2016, theorem 2.5).

5. Monte Carlo Experiment

This section uses a Monte Carlo experiment to inspect the finite sample performance of the classification and estimation procedure introduced in section 4 when applied to group-specific impulse responses in the presence of latent group heterogeneity.

5.1. Data generating process

Each Monte Carlo sample consist of consists of a panel data $\{(y_{i,t}, \mathbf{x}_t)\}$ for $i = 1, \ldots, N$ and $t = 1, \ldots, T$ that is generated according to,

$$y_{i,t} = \sum_{h=0}^{12} x_{t-h}\beta_{i,h} + \varepsilon_{i,t} \tag{13}$$

where x_t is an aggregate shock such that $x_t \sim \mathcal{N}(0, 1)$ and i.i.d. across t, the idiosyncratic shocks $\varepsilon_{i,t} \sim \mathcal{N}(0, 1)$ are i.i.d. across i and t and x_t and $\varepsilon_{i,t}$ are mutually independent. There are two latent groups ($K_0 = 2$) and the group-specific impulse responses are parametrised using a Gaussian basis function as in Barnichon and Mathes (2018). In particular,

$$\beta_{i,h} = \begin{cases} 0.15 \times \exp\left\{-\left(\frac{h-4}{25}\right)^2\right\}, \text{if } i \in G_1 \\ -0.15 \times \exp\left\{-\left(\frac{h-4}{25}\right)^2\right\}, \text{if } i \in G_2 \end{cases} \tag{14}$$

which results in symmetric impulse responses for groups 1 and 2 as depicted in figure 1. Individuals are assigned to group 1 if they are indexed by an odd number and assigned to group 2 if they are indexed by an even number so that, for each sample generated, half of the individuals belong to each group. Sample sizes of size $N = \{100, 200\}$ and time spans $T = \{40, 80\}$ are considered.[8] For each possible combination of N and T, 250 Monte Carlo samples are generated.

5.2. Estimation and Classification

For each Monte Carlo sample generated two alternative ways of estimating the impulse responses are considered. The first one is by focusing directly on the

[8] Even a value of $T = 80$ is still relatively small compared to what is typically used in the literature using local projections to estimate impulse responses. As documented by Herbst and Johannsen (2022) the median value for T across the 100 "most relevant" papers citing Jordà (2005) in Google scholar is around 95.

moving average representation and, in that case, the C-Lasso objective function is given by (11) with $\mathbf{x}'_{i,t} = [x_t, x_{t-1}, \ldots, x_{t-H}]$. The second way of estimating impulse responses is trough the use of local projections (Jordà, 2005) which use a sequence of regressions of $y_{i,t}$ on $x_{i,t-h}$ to estimate the impulse response for individual i at horizon h. The case of local projections can be accommodated in the C-Lasso framework described in section 4 by replacing (11) by the following modified C-Lasso objective function,

$$\widetilde{Q}^{(K)}_{NTH,\lambda_1}(\mathbf{b}, \mathbf{a}) = \frac{1}{N(H+1)T} \sum_{i=1}^{N} \sum_{h=0}^{H} \sum_{t=1}^{T} \left(y_{i,t} - x_{i,t-h} \mathrm{b}_{h+1,i}\right)^2 + \frac{\tilde{\lambda}_1}{N} \sum_{i=1}^{N} \prod_{k=1}^{K} \|\mathbf{b}_i - \mathbf{a}_k\| \tag{15}$$

where $\mathrm{b}_{h+1,i}$ is the $(h+1, i)$-*th* element of the matrix $\mathbf{b}$ and $\tilde{\lambda}_1$ is a tunning parameter that tends to zero as $(N, T) \to \infty$.[9] The local projections C-Lasso estimates of $\boldsymbol{\beta}$ and $\boldsymbol{\alpha}$ are obtained by minimising (15) with respect to $\mathbf{b}$ and $\mathbf{a}$. Given these estimates the individual classification, determination of number of groups and post-Lasso estimates are obtained in an analogous way as described in sections 4.1 to 4.3, except the tunning parameters that are adjusted to reflect the effective number of observations per cross-sectional unit that is different for the local projections case.

Estimation of impulse responses through local projections has become increasingly popular over the last decade. Among the advantages of local projections cited in Jordà (2005) are their flexibility and the fact that they are more robust to misspecification of the moving average representation if it arises from the inversion of a misspecified vector autoregression. Notice, however, that in the present

[9] The modification needed to estimate impulse responses through local projections using the C-Lasso framework is more easily seen in matrix form. Let $\mathbf{X}_i$ be defined as in assumption 1 and $\mathbf{y}_i = [y_{i,1}, \ldots, y_{i,T}]'$. The first term on the right-hand-side of (11) is given by $\frac{1}{NT}\sum_{i=1}^{N}(\mathbf{y}_i - \mathbf{X}_i\mathbf{b}_i)'(\mathbf{y}_i - \mathbf{X}_i\mathbf{b}_i)$. Estimation through local projections simply requires replacing this term by $\frac{1}{N\tilde{T}}\sum_{i=1}^{N}(\tilde{\mathbf{y}}_i - \tilde{\mathbf{X}}_i\mathbf{b}_i)'(\tilde{\mathbf{y}}_i - \tilde{\mathbf{X}}_i\mathbf{b}_i)$ where $\tilde{T} = (H+1)T$, $\tilde{\mathbf{y}}_i = \mathbf{1}_{N\times 1} \otimes \mathbf{y}_i$ and $\tilde{\mathbf{X}}_i = \oplus_{j=1}^{H+1}\mathbf{x}_{*,j}$ and $\mathbf{x}_{*,j}$ denotes the *j-th* column of $\mathbf{X}_i$.

Monte Carlo experiment the moving average representation is *not* misspecified and, hence, it is not expected that the impulse responses estimated from local projections to display better statistical properties than those estimated directly from the moving average representation. The purpose of including impulse response estimation through local projections in the present exercise is simply to illustrate how they can be accommodated by the C-Lasso framework.

Tuning parameters. Determination of the number of groups and individual classification requires the researcher needs to specify the tuning parameters λ_1 in (11) and ρ_{NT} in (12). The assumptions on λ_1 and ρ_{NT} needed to derive the asymptotic properties highlighted in section 4.4 are satisfied for any λ_1 such that $\lambda_1 \propto T^{-a}$ for any $a \in (0, -1/2)$ and any ρ_{NT} that can be written as $\rho_{NT} \propto (NT)^{-b}$ for any $b \in (0, 1)$. Even though asymptotically the choice of the tuning parameters is irrelevant as long as they satisfy these conditions, in finite samples their choice can be crucial. In this Monte Carlo experiment the values of the tuning parameters used to estimate impulse responses through the moving average representation are $\lambda_1 = c_\lambda s_Y^2 T^{-1/3}$ and $\rho_{NT} = \frac{2}{3}(NT)^{-\frac{2}{3}}$ where s_Y^2 denotes the sample variance of $y_{i,t}$ and c_λ is set equal to 2.[10,11] For the estimation via local projections the tuning parameters are adjusted to reflect the effective number of observations per cross-sectional unit, that is, $\tilde{\lambda}_1 = \tilde{c}_\lambda s_Y^2 \tilde{T}^{-1/3}$ and $\tilde{\rho}_{NT} = \frac{2}{3}(N\tilde{T})^{-\frac{2}{3}}$ where $\tilde{c}_\lambda$ is equal to 2 and $\tilde{T} = (H + 1)T$.

[10] Su, Shi and Phillips (2016) use $c_\lambda \in [0.125, 0.25, 0.5, 1, 2]$ and select the value of c_λ *jointly* with the number of latent groups k to minimise the information criterion in (12). For a small number of Monte Carlo replications it was found that jointly determining c_λ with the number of groups did not had effect on the results whilst substantially increasing the computational costs. For this reason or each Monte Carlo replication c_λ is kept fixed equal to 2. In the empirical application in section 6, the value of c_λ is grid search over the same grid used in Su, Shi and Phillips (2016) and jointly chosen with the number of groups to minimize the information criterion.

[11] For linear models Su, Shi and Phillips (2016) use $\rho_{NT} = \frac{2}{3}(NT)^{-\frac{1}{2}}$. In numerical experiments, for the DGP here considered I found this value for the tuning parameter tends to over-select number of groups that is *smaller* than the true number of groups. I have experimented for values $\rho_{NT} = c_1(NT)^{-c_2}$ for $c_1, c_2 \in (0, 1)$ and found that $\rho_{NT} = \frac{2}{3}(NT)^{-\frac{2}{3}}$ tends to select the correct number of groups with a higher frequency.

5.3. Monte Carlo Results

The results from the Monte Carlo experiment are reported in tables 1 to 3 and in figures 2 and 3. They can be summarised as follows:

Determination of the number of latent groups. For each DGP considered, table 1 shows the frequency that different number of groups is chosen across Monte Carlo replications. When the number of groups is based on the estimation of the moving average representation the IC-based group determination procedure always identifies the correct number of latent groups except in 1% percent of the samples for $N = 100$ and $T = 40$ where the IC picks one latent group. When the number of groups is based on the estimation of local projections the performance of the group determination procedure deteriorates, specially for the two DGPs with $T = 40$ where the IC selects one latent group more often than two latent groups. As expected, as N and T increase the frequency that the true number of latent groups increases and, in particular, for $N = 200$ and $T = 80$ the IC selects the correct number of latent groups in 82% of the Monte Carlo samples.

Individual classification. The average individual misclassification rates across Monte Carlo replications for each DGP is reported in figures 2 and 3. This figure consists of the share of individuals that are assigned to a group they do *not* belong averaged across Monte Carlo samples. From figure 2, when impulse responses are estimated through the moving average representation this figure is under 4% for $T = 40$ and under 1% for $T = 80$ which is suggestive that the classifier proposed in 4.1 tends to classify individuals to the correct group in finite samples too. From figure 3, when estimating impulse responses through local projections are of the order of 30% for $T = 40$% and of the order of 10% when $T = 80$. This inferior performance for the local projections case is justified by the fact that coefficient estimates from local projection regressions have much higher sampling variability since the error term in those regressions includes not only the original error term from the moving average representation ($\varepsilon_{i,t}$) but also all the other leads and lags that are not included.[12] With higher sampling variability that estimated individual impulse responses for individuals from group 1 (group 2) end up being closer to the impulse response from group 2 (group 1) which when applying the classifier leads that individual to be assigned to the wrong group.

Post-Lasso impulse responses. The estimated post-Lasso impulse responses are plotted against the true group impulse responses in figures 2 and 3. For the case where impulse responses are estimated directly through the moving average representation (figure 2) the estimated impulse responses almost overlap the true impulse responses which indicates the absence of bias. For the case where $T = 40$ there is some small discrepancies that can be justified by the slightly higher misclassification rate than for the $T = 80$ case. For the case of impulse responses estimated via local projections (figure 3), there is some bias specially for $T = 40$ where the misclassification rate is of the order of 30%, however, and as expected from (9), the bias substantially decreases for the $(N, T) = (200, 80)$ case when the average misclassification rate drops to 7%. Moreover, the difference between the 90th and the 10th percentiles of the sampling distribution is always smaller in figure 2 than in 3 which echoes the fact that the sampling variance of the impulse responses estimated via local projections is higher than those estimated directly through the moving average representation (see footnote [12]). The analysis of figures 2 and 3 is complemented with the figures in tables 2 and 3 that compares the bias, variance and mean squared error for the impulse response estimates for the two groups at horizon $h = 4$ (*i.e.* the peak of the impulse responses). The figures

[12] Notice that the true data generating process is given by $y_{i,t} = \sum_{h=0}^{12} x_{t-h}\beta_{i,h} + \varepsilon_{i,t}$. For a given horizon h, the local projection regression is given by $y_{i,t} = x_{t-h}\beta_{i,h} + \tilde{\varepsilon}_{i,t}$ where $\tilde{\varepsilon}_{i,t} = \varepsilon_{i,t} + \sum_{j \neq h}^{12} x_{t-j}\beta_{i,j}$. Since the aggregate shocks are mean zero and iid, the fact that they are omitted does *not* cause bias or inconsistency in the local projection estimates of impulse responses but it *does* increase their sampling variance vis-a-vis the estimates obtained through the estimation of the moving average representation.

in both tables numerically illustrate the theoretical results from proposition 1. The full heterogeneity estimator has essentially no bias but a higher variance than the post-Lasso estimator, whilst the post-Lasso estimator has some bias, since in finite samples it does not achieve perfect classification of individuals into groups, but a smaller variance. Most importantly, in MSE terms the post-Lasso estimator is always preferred to the full heterogeneity one. In particular, for the estimation through the moving average representation the decrease in MSE of the post-Lasso estimator vis-à-vis the full heterogeneity over one full order of magnitude. For the estimation through local projection the MSE of the post-Lasso estimator is one to two thirds smaller than the MSE of the full heterogeneity estimator.

6. Aggregate IST shocks and Firm level debt revisited

This section uses the C-Lasso classification and estimation procedure introduced in section 4 to revisit the estimation of impulse responses of firm-level debt to an aggregate investment specific shock originally studied by Drechsel (2023).

6.1. Background

Motivated by microeconomic evidence on corporate borrowing in the US that unveals a direct connection between firms' current earnings and their access to debt, Drechsel (2023) studies the macroeconomic implications of the so-called *earnings-based borrowing constraints*. First, in a prototype business cycle model the author shows that depending on the type of borrowing constraint used firm-level debt responds differently to a permanent investment shock. More precisely, in a setting where firms face a standard collateral constraint their debt decreases in response to a positive investment investment shock whereas if they face an earnings constraint their debt increases following that same shock (see Drechsel, 2023, figure 2).

To empirically test this model prediction, the author uses a time-series of investment specific technology (IST) shocks identified from an SVAR using long-run restrictions combined with a firm-level panel containing firm characteristics and information on the types of covenants included in their debt contracts.[13] The baseline specification to test the model predictions is given by,

$$\begin{aligned}\log\left(b_{i,t+h}\right) = {} & \alpha_h + \beta_h \hat{u}_{IST,t} + \gamma \boldsymbol{X_{i,t-1}} \\ & + \beta_h^{earn} \mathbb{1}_{i,t,earn} \times \hat{u}_{IST,t} + \alpha_h^{earn} \mathbb{1}_{i,t,earn} \\ & + \beta_h^{coll} \mathbb{1}_{i,t,earn} \times \hat{u}_{IST,t} + \alpha_h^{coll} \mathbb{1}_{i,t,coll} + \delta t + \eta_{i,t+h}\end{aligned} \tag{16}$$

where $b_{i,t}$ is the quarterly level of firms' debt liabilities, $\mathbb{1}_{i,t,earn}$ and $\mathbb{1}_{i,t,coll}$ are dummy variables that capture whether the firm is subject to earnings-related covenants or uses collateral, $\hat{u}_{IST,t}$ is the IST shock identified based on long-run restrictions, t is a linear time trend and $\boldsymbol{X_{i,t-1}}$ is vector of controls that includes $\log\left(b_{i,t-1}\right)$, 3-digit industry-level fixed effects, firm size, firm-level real sales growth and a variable constructed from SVAR residuals that is meant to capture macroeconomic shocks other than investment shocks.[14] At a given horizon h, the impulse response of an "earnings-based borrower" ("collateral-based borrower") is given by the sum of the coefficients $\beta_h + \beta_h^{earn}$ ($\beta_h + \beta_h^{coll}$) and, hence, in terms of regression coefficients the model predictions to be tested are $\beta_h + \beta_h^{earn} > 0$ and $\beta_h + \beta_h^{coll} < 0$.

The results based on (16) are presented in Drechsel (2023, figure 7) and are largely in line with the model implied impulse responses. In the data, across a wide range of specifications, debt of earnings-based borrowers reacts positively to an IST shock whereas debt for collateral-based borrowers declines in response to that same shock.

[13] This panel is obtained by merging the Dealscan dataset with Compustat data. For more details on the construction of this dataset and which variables are used in which specification refer to section 4.3.2 in Drechsel (2023).

[14] More specifically, $\mathbb{1}_{i,t,earn}$ is equal to 1 if a given firm issues a loan with at least one earnings covenant and $\mathbb{1}_{i,t,earn}$ is equal to one if the debt issued by the firm is secured by specific assets.

6.2. Revisiting the responses of debt to IST shock

Following the bulk of the existing literature that has conducted heterogeneity analysis on impulse response functions, Drechsel (2023) adopts the *ex-ante classification approach* described in section 3. Following the analysis of micro data on firm-level debt issuances and theoretically grounded by model predictions, the author groups firms depending on whether they are "earnings-based borrowers" or "collateral-based borrowers" and estimates the impulse responses to an IST shocks for each of these two groups of firms based on (16). In light of proposition 1, this approach can be subject a misclassification bias if the ex-ante grouping of firms, based on whether their borrow against collateral or earnings, does not coincide with the true grouping that underlies the heterogeneity firm-level impulse responses observed in the data. To investigate whether this is the case this section re-estimates the impulse responses of firm level debt to an IST shock using the C-Lasso classification and estimation procedure proposed in section 4. In order to do so, the following moving average version of (16) is considered,

$$\tilde{b}_{i,t} = \sum_{h=0}^{H} \beta_{i,h} \tilde{\hat{u}}_{IST,t} + \epsilon_{i,t} \tag{17}$$

where $\tilde{b}_{i,t}$ denotes the residuals of a regression of $\log(b_{i,t})$ on a constant and a linear time trend and, similarly, $\tilde{\hat{u}}_{IST,t}$ denotes the residuals of $\hat{u}_{IST,t}$ on a constant and a linear time trend. The identified IST shock is assumed to be strictly exogenous. Since all the theoretical results for the C-Lasso are derived for a balanced panel and all the empirical applications in Su, Shi and Phillips (2016) are also focused on balanced panels, specification (17) is estimated using the C-Lasso approach based on a balanced version of the Drechsel (2023) dataset. The final balanced panel contains 746 firms and 76 quarters spanning the period from 1997Q1 to 2015Q4.[15]

[15] To reach this balanced panel from the original Drechsel (2023) dataset, I first exclude the 12 periods of the dataset that are lost due to the lagging of the IST shock then, on the resulting sample, I remove all the firms that have missing values of debt for any quarter between 1997Q1 to 2015Q4.

6.3. Results

Number of latent groups and post-Lasso IRFs. As illustrated in figure 4, the IC-based group determination procedure identifies two-latent groups. The post-Lasso IRF estimates for each of these groups along with firm-specific IRFs and the IRF estimated by pooling all the firms are depicted in figure 5. In line with the theoretical predictions in Drechsel (2023), one of the two latent groups responds *positively* to an IST shocks whilst the second group responds *negatively*. In addition, two points are also worth noting from figure 5. First, there is significant heterogeneity among individual-specific impulse responses and without any ex-ante theoretical reason to group individuals it would be difficult to identify the group pattern identified by the C-Lasso just by looking at the cloud formed by individual impulse responses. Second, despite the differences in the specifications (16) and (17) and the sample composition, it is reassuring to see that the pooled IRF in figure 5 has the same shape as the pooled impulse responses depicted in figure 6 in Drechsel (2023), that, the response of debt is increasing up until 2 years after the shock and then starts declining.[16]

Firm characteristics across the two groups. Despite the two latent groups and their associated responses being in line with predictions from Drechsel (2023), the most important aspect to be tested is whether indeed the group that responds positively to an IST shock contains a disproportionately higher share of earnings-based borrowers vis-à-vis the group that responds negatively. In total there are 746 firms, 224 of which where classified into the positive response group whereas the remaining 522 where classified into the negative response group. Table 4 looks at four different firm characteristics across these two groups. Panel A looks at the

[16] Despite the similar shape of the IRFs there are some differences in terms of magnitudes. The on impact response from the pooled specification in Drechsel (2023) is zero and the peak of the response, that occurs 7 quarters after the shock, is around 2.5%. In figure 5, the on impact response of debt is roughly -2% whereas the peak response is almost 6%.

relative proportions of earnings and collateral borrowers across the two groups. To compute these proportions a firm is classified as earnings-based borrower if over the whole sample it has more debt earnings based debt issuances and it is classified as a collateral-based borrower if over the whole sample it has more collateral debt issuances than earnings based ones.[17] The majority of firms in both groups are earnings-based borrowers, however, in line in the theoretical predictions in Drechsel (2023) the share of earnings-borrowers is larger in the group that responds positively to an IST shock although the difference is not statistically significant. Panel B looks at the two measures of the share of the intangible assets as a share of total assets. Theoretically, one would expect firms that have larger share of intangible assets to borrow more against earnings since intangible assets cannot be used as collateral and, hence, the larger the share of intangible assets the more likely the firm is to respond positively to an IST shock. This is indeed the case, as the average of both measures of intangibility are higher for group 1, yet the difference is quantitatively small and not statistically significant. Panel C looks at three different measures of firm size as in Dang, Li and Yang (2018). In theory, one would expect smaller and younger firms to have less collateral to pledge and, hence, to borrow more against their future earnings. Again this is confirmed in the data, since for the three measures considered the average firm size is smaller in the group 1 and the difference is statistically significant at the 10% level when size is measured by firm's total assets. Finally, panel D looks at sectorial composition of each of the two groups. In this respect, group 1 has a statistically significant higher share of firms in the materials and industrial sectors whereas group 2 has a significantly higher share of firms in the consumer staples and utilities sectors.

[17] Notice this criteria is slightly different than the one used in Drechsel (2023) since the dummies $\mathbb{1}_{i,t,earn}$ and $\mathbb{1}_{i,t,coll}$ are only defined for quarters where a debt issuance for firm i appears in the Dealscan dataset. These dummies are defined relative to a specific debt issuance and, hence, they can vary over time (*e.g.* a given firm can very well issue debt against collateral in a given date and issue another debt contract with earnings covenants in other date. To compute the shares reported in Panel A of table 4 a firm is classified as an earnings-based borrower if $\sum_{t=1}^{T} \mathbb{1}_{i,t,earn} > \sum_{t=1}^{T} \mathbb{1}_{i,t,coll}$ and classified as a collateral-based borrower if $\sum_{t=1}^{T} \mathbb{1}_{i,t,earn} < \sum_{t=1}^{T} \mathbb{1}_{i,t,coll}$. Otherwise, if $\sum_{t=1}^{T} \mathbb{1}_{i,t,earn} = \sum_{t=1}^{T} \mathbb{1}_{i,t,coll}$ the firm is not classified.

Explaining group membership. To estimate the impact of each of the variables in table 4 on the probability that a given firm belongs to each of the two groups, a Logit model is estimated using as dependent variable a dummy that is equal to 1 if the firm belongs to group 1. The average marginal effects of each variable across different specifications are reported in table 5. The results in this table when including each group of explanatory variables at a time (columns 1 to 4) largely corroborate the conclusions based on the analysis of group characteristics in table 4. The specification in column 5 includes all the covariates at the same time. In this specification the effects of the first three variables are quantitatively small and not statistically significant. The only two statistically significant covariates are the dummies for the consumer staples and utilities sectors. In particular, a firm that is classified as consumer staples (utilities) is, on average, 40 p.p. (21 p.p.) *less* likely to belong to the group that for which debt increases following an IST shock. In summary, the group that responds positively to an IST shock is composed by firms that: (*i*) that are relatively smaller; (*ii*) have higher share of *intangible assets*, (*iii*) tend to be *earnings-based borrowers* and (*iv*) do *not* belong to the *consumer staples* or *utilities* sectors.[18] This findings are largely in line with the theoretical predictions in Drechsel (2023), but also suggest that, on top of whether a firm tends to borrow against earnings or collateral, the specific *sector* that the firm operates also plays a role in determining whether it will respond positively or negatively to an aggregate IST shock.[19]

[18] The type of exercise explaining group membership is in spirit to a principal component analysis where the data alone selects the orthogonal factors that explain the correlations observed in the data and ex-post the researcher searches for appropriate names for these factors. In the present context, the C-Lasso approach selects the number of groups, the individual classification and the group-specific IRFs based solely on the data and it is based on the analysis of individual characteristics across characteristics that a name for each group is determined.

[19] In addition to the specification in (17), an alternative specification based on local projections is also estimated. In that specification, the IC-based group determination procedure identifies only one latent group. Given the superior performance of the moving average representation in identifying the correct number of groups in the Monte Carlo experiment, the discussion in the main text focuses on the moving average representation. Nonetheless, the estimates of the local projection specification *conditional* on two latent groups yields post-Lasso IRFs that are similar to the post-Lasso IRFs depicted in figure 5. Moreover, the group classification based on the local projection specification has an overlap of 91% with the individual classification based on the moving average

7. Conclusion

This paper studied the estimation of heterogeneous impulse responses in the presence of latent group heterogeneity. It showed that the common approach in the literature based on the ex-ante grouping of individuals according to some external criteria or observable explanatory variables can lead to misleading conclusions. More precisely, the choice between an estimator of group-specific impulse responses based on an ex-ante grouping of individuals and estimating individual-specific impulse responses entails a bias-variance tradeoff. Motivated by this tradeoff, this paper proposed an alternative methodology based on the C-Lasso to estimate group-specific impulse responses. A Monte Carlo experiment demonstrated good finite-sample performance of this methodology both in classification of individuals into different groups and estimation of group specific impulse responses. An application of this methodology to study firm level debt responses to an aggregate IST shock based on Drechsel (2023) identified two latent groups. One group of firms for which firm-level debt responds positively to an IST shock and other for which the response is negative. The group of firms for which debt increases in response to a positive IST shock is composed by firms that are relatively smaller, have higher share of *intangible assets*, tend to be *earnings-based borrowers* and do *not* belong to the *consumer staples* or *utilities* sectors.

I conclude by highlighting two dimensions along which the results from the present paper can be extended and applied. On the methodological side, the methodology proposed could be extended to include the possibility that the shocks are used as instruments and a more thorough Monte Carlo study could be conducted to search for values of the fine tuning parameter that improve finite-sample performance in terms of group determination when using local projections. On the applications front, the methodology here introduced could be used to either

representation.

test alternative transmission channels from aggregate shocks to the cross-section (in the spirit of the application in section 6) or to identify the most important dimensions that drive heterogeneous responses to aggregate shocks and use that information to inform the theoretical modeling of DSGE models featuring heterogeneity across firms and/or households.

References

Baltagi, Badi H., *Econometrics*, 4th ed., Springer, Berlin, Heidelberg, 2008.

Barnichon, Regis and Christian Mathes, "Functional Approximation of Impulse Responses," *Journal of Monetary Economics*, 2018, *99*, 41–55.

Bernanke, Ben S. and Mark Gertler, "Inside the Black Box: The Credit Channel of Monetary Policy Transmission," *Journal of Economic Perspectives*, 1995, *9* (4), 27–48.

Bilbiie, Florin O., "Limited asset markets participation, monetary policy and (inverted) aggregate demand logic," *Journal of Economic Theory*, 2008, *140*, 162–196.

Bonhomme, Stéphane and Elena Manresa, "Grouped Patterns of Heterogeneity in Panel Data," *Econometrica*, 2015, *83* (3), 1147–1184.

Browning, Martin and Jesus M. Carro, "Dynamic binary outcome models with maximal heterogeneity," *Journal of Econometrics*, 2014, *178* (2), 805–823.

Carlino, Gerald and Robert Defina, "The Differential Regional Effects of Monetary Policy: Evidence from the U.S. States," *Journal of Regional Science*, 1999, *39* (2), 339–358.

Carvalho, Carlos, "Heterogeneity in Price Stickiness and the Real Effects of Monetary Shocks," *Frontiers in Macroeconomics*, December 2006, *2* (1).

Clayton, Christopher, Xavier Jaravel, and Andreas Schaab, "Heterogeneous Price Rigidities and Monetary Policy," August 2018.

Cloyne, James, Clodomiro Ferreira, and Paolo Surico, "Monetary Policy when Households have Debt: New Evidence on the Transmission Mechanism," *Review of Economic Studies*, January 2019, *87*, 102–129.

_ , _ , Maren Froemel, and Paolo Surico, "Monetary Policy, Corporate Finance and Investment," *Journal of European Economic Association*, March 2023.

Coibion, Olivier, Yuriy Gorodnichenko, Lorenz Kueng, and John Silvia, "Innocent Bystanders? Monetary Policy and Inequality," *Journal of Monetary Economics*, June 2017, *88*, 70–89.

Cravino, Javier, Ting Lan, and Andrei A. Levchenko, "Price stickiness along the income distribution and the effects of monetary policy," *Journal of Monetary Economics*, 2020, *110*, 19–32.

Dang, Chongyu, Zhichuan Li, and Chen Yang, "Measuring firm size in empirical corporate finance," *Journal of Banking and Finance*, 2018, *86*, 159–176.

Drechsel, Thomas, "Earnings-Based Borrowing Constraints and Macroeconomic Fluctuations," *American Economic Journal: Macroeconomics*, April 2023, *15* (2), 1–34.

Frisch, Ragnar, "Propagation problems and impulse problems in dynamic economics," in "Economic Essays in Honor of Gustav Cassel," London: Allen & Unwin, 1933, pp. 171–205.

Galí, Jordi, J. David López-Salido, and Javier Vallés, "Understanding the Effects of Government Spending on Consumption," *Journal of the European Economic Association*, March 2007, *5* (1), 227–279.

Gertler, Mark and Peter Karadi, "Monetary policy surprises, credit costs, and economic activity," *American Economic Journal: Macroeconomics*, 2015, *7* (1), 44–76.

_ **and Simon Gilchrist**, "Monetary Policy, Business Cycles, and the Behavior of Small Manufacturing Firms," *Quarterly Journal of Economics*, May 1994, *109* (2), 309–340.

Greene, William H., *Econometric Analysis*, 5th ed., Upper Saddle River, New Jersey 07458: Prentice Hall, 2003.

Hayashi, Fumio, *Econometrics*, first ed., Princeton University Press, 2000.

Herbst, Edward P. and Benjamin K. Johannsen, "Bias in Local Projections," November 2022.

Hsiao, Cheng, *Analysis of Panel Data*, Vol. Econometric Society Monographs, Cambridge University Press, 2014.

Jeenas, Priit, "Firm balance sheet liquidity, monetary policy shocks, and investment dynamics," June 2019.

Jordà, Òscar, "Estimation and Inference of Impulse Responses by Local Projections," *American Economic Review*, March 2005, *95* (1), 161–182.

Kasahara, Hiroyuki and Katsumi Shimotsu, "Nonparametric identification of finite mixture models of dynamic discrete choices," *Econometrica*, 2009, *77* (1), 135–175.

Kashyap, A., O. Lamont, and J. Stein, "Credit conditions and the cyclical behavior of inventories," *Quarterly Journal of Economics*, 1994, *109* (3), 565–592.

Lewis, Daniel, Davide Melcangi, and Laura Pilossoph, "Latent Heterogeneity in the Marginal Propensity to Consume," September 2022.

Lin, Chang-Ching and Serena Ng, "Estimation of Panel Data Models with Parameter Heterogeneity when Group Membership is Unknown," *Journal of Econometric Methods*, 2012, *1* (1).

Lu, Xun and Liangjun Su, "Determining the number of groups in latent panel structures with an application to income and democracy," *Quantitative Economics*, 2017, *8*, 729–760.

Ottonello, Pablo and Thomas Winberry, "Financial Heterogeneity and the Investment Channel of Monetary Policy," *Econometrica*, November 2020, *88* (6), 2473–2502.

Ramey, Valerie A., "Macroeconomic Shocks and Their Propagation," in "Handbook of Macroeconomics," Vol. 2, Elsevier, 2016, chapter 2.

Romer, Christina and David Romer, "A New Measure of Monetary Shocks: Derivation and Implications," *American Economic Review*, 2004, *94* (4), 1055–1084.

Sarafidis, Vasilis and Neville Weber, "A Partially Heterogeneous Framework for Analyzing Panel Data," *Oxford Bulletin of Economics and Statistics*, April 2015, *77* (2), 274–296.

Slutzky, Eugen, "The Summation of Random Causes as the Source of Cyclic Processes," *Econometrica*, April 1937, *5* (2), 105–146.

Su, Liangjun, Zhentao Shi, and Peter C. B. Phillips, "Identifying Latent Structures in Panel Data," December 2014. Cowles Foundation Discussion Paper No. 1965.

_ , _ , and _ , "Identifying Latent Structures in Panel Data," *Econometrica*, November 2016, *84* (6), 2215–2264.

Sun, Yixiao, "Estimation and Inference in Panel Structure Models," October 2005. Working Paper, Dept. of Working Paper, Dept. of Economics, UCSD.

Wang, Wuyi, Peter C. B. Phillips, and Liangjun Su, "Homogeneity Pursuit in Panel Data Models: Theory and Application," *Journal of Applied Econometrics*, September 2018, *33* (6), 797–815.

Wong, Arlene, "Refinancing and The Transmission of Monetary Policy to Consumption," May 2019.

Figures and Tables

Figure 1: Heterogeneous impulse responses under latent group heterogeneity

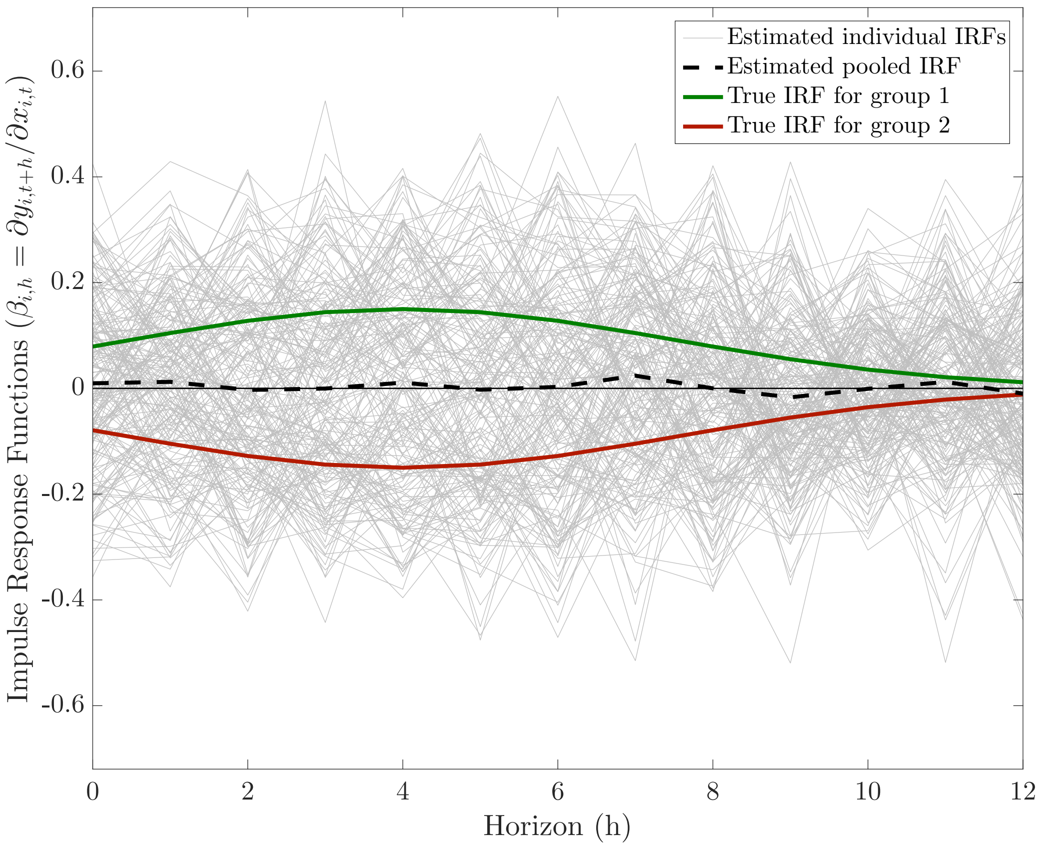

The plot is generated based on artificially generated panel data set with 200 individuals and 40 time periods. Half of the individuals have their true impulse responses given by the green line and the other half by the red line. The grey lines are individual-specific estimated impulse responses. The dashed-black line is the estimated impulse response obtained by pooling all the individuals and ignoring coefficient heterogeneity. The data generating process is described in section 5.1.

Figure 2: Post Lasso IRF estimates based on moving average representation

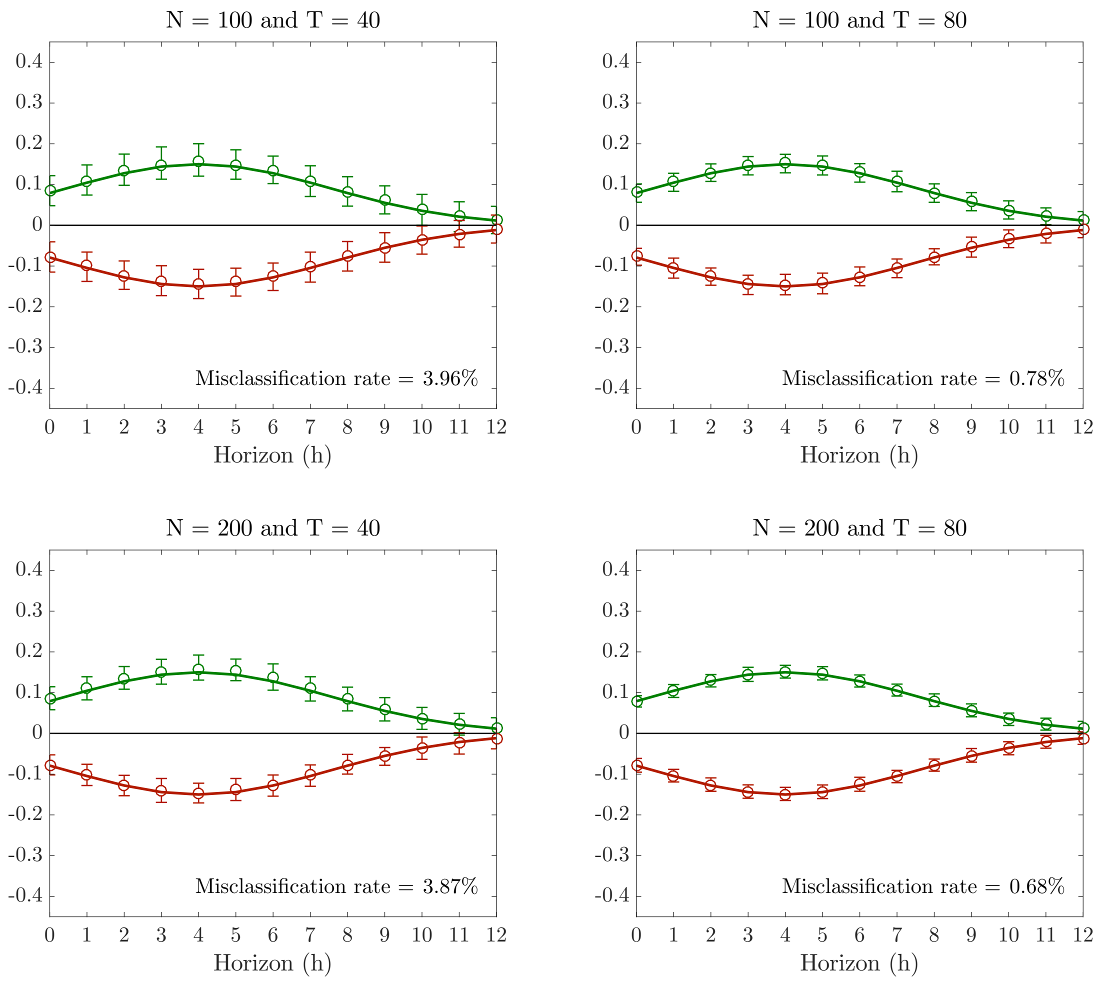

The solid lines are the true impulse responses for each group as defined in (14). The hollow circles represent the means of the sampling distribution of post-Lasso impulse response computed across Monte Carlo replications. The vertical line contains the interval from the 10th to 90th percentile of the post-Lasso impulse response estimates computed across Monte Carlo replications. For each Monte Carlo sample the misclassification rate is computed as $\frac{1}{N}\sum_{i=1}^{N}(\mathbb{1}\{i \in G_1 \cap \widehat{G}_2\} + \mathbb{1}\{i \in G_2 \cap \widehat{G}_1\})$ and the misclassification rates reported are the average misclassification rate across Monte Carlo samples.

Figure 3: Post Lasso IRF estimates based on local projections

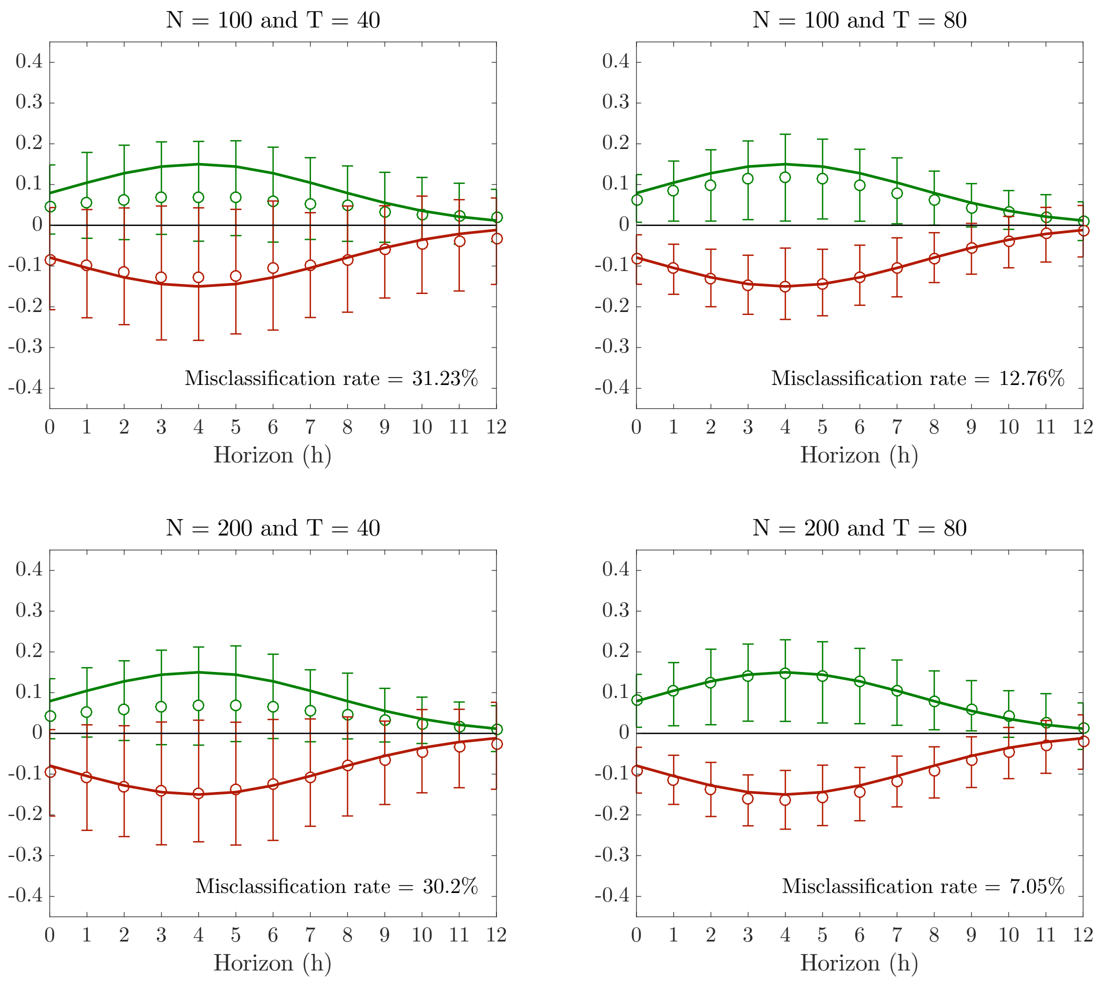

The solid lines are the true impulse responses for each group as defined in (14). The hollow circles represent the means of the sampling distribution of post-Lasso impulse response computed across Monte Carlo replications. The vertical line contains the interval from the 10th to 90th percentile of the post-Lasso impulse response estimates computed across Monte Carlo replications. For each Monte Carlo sample the misclassification rate is computed as $\frac{1}{N}\sum_{i=1}^{N}(\mathbb{1}\{i \in G_1 \cap \widehat{G}_2\} + \mathbb{1}\{i \in G_2 \cap \widehat{G}_1\})$ and the misclassification rates reported are the average misclassification rate across Monte Carlo samples.

Figure 4: Group determination for IST shock x firm-level debt responses

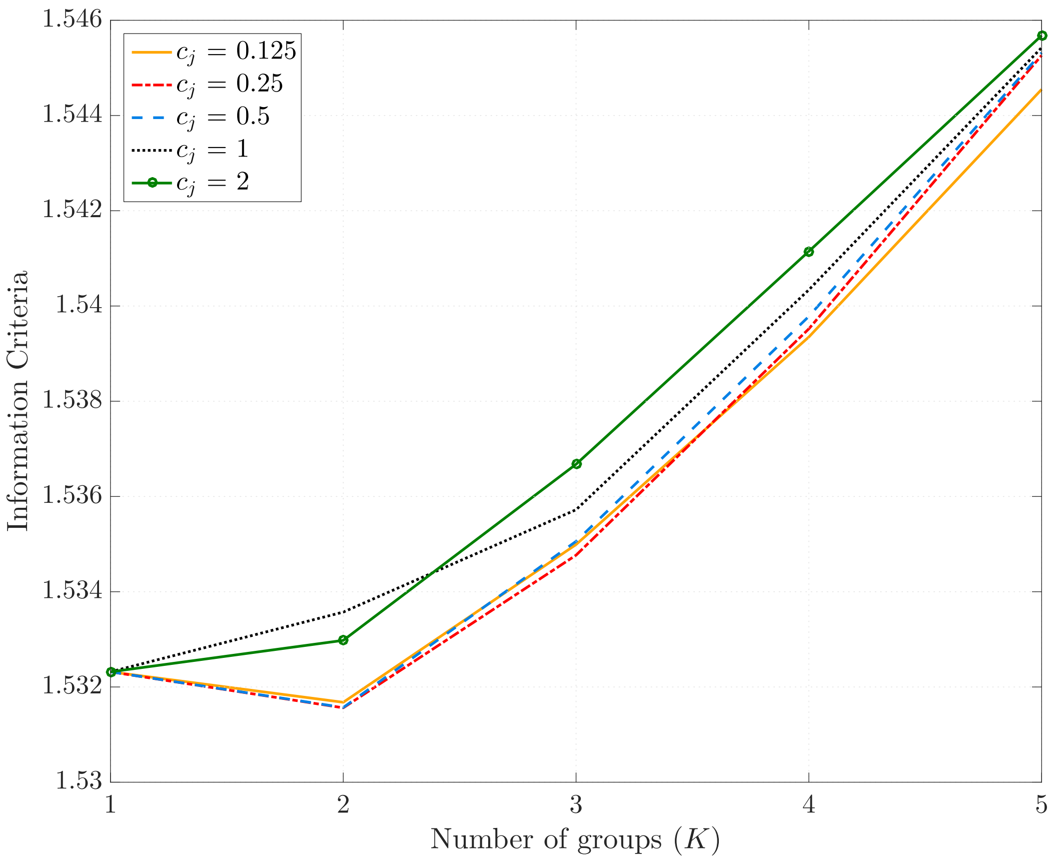

Each line reports values of IC(K, λ_1), as defined in (12), computed from the Drechsel (2023) dataset for alternative values of K and for $\lambda_1 = c_j s_Y^2 T^{\frac{1}{3}}$. The combination of (K, c_j) that minimises the IC is given by $(2, 0.25)$.

Figure 5: Post Lasso IRFs for responses of firm-level debt to IST shock

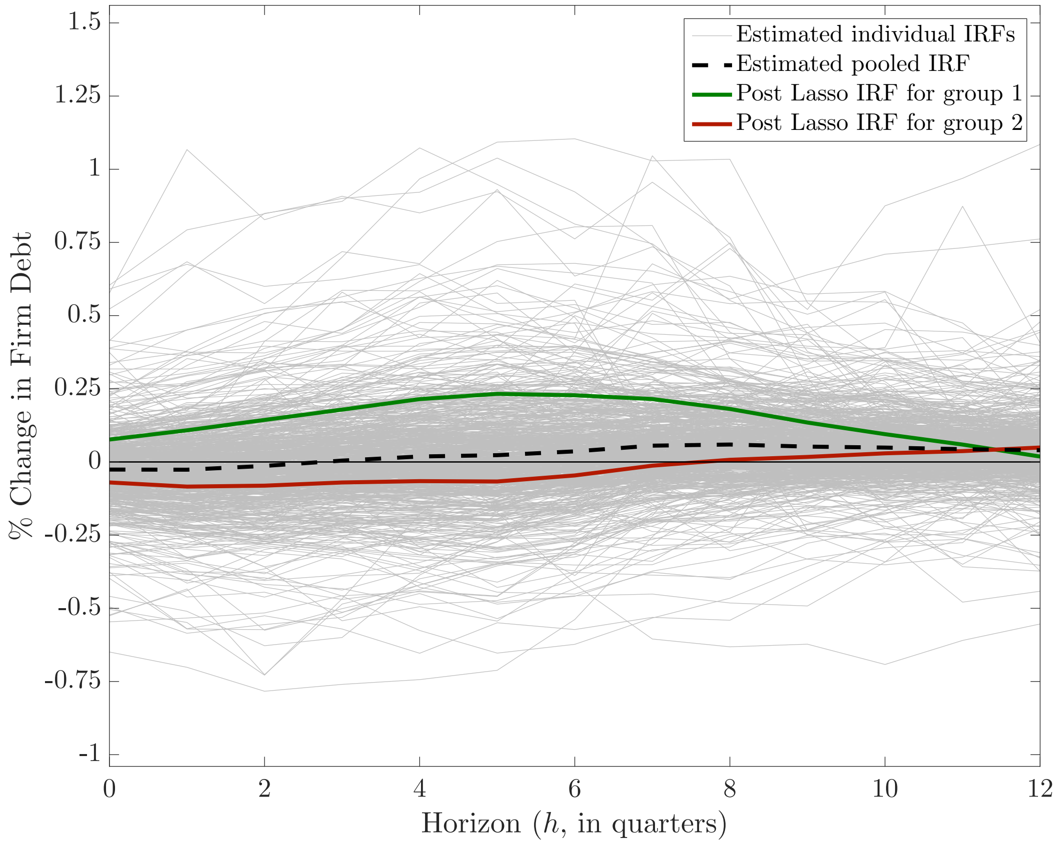

The solid green and red lines are the estimated post-Lasso impulse responses for the two latent groups identified in the Drechsel (2023) dataset. The dashed black line plot the impulse response obtained by pooling all the firms together. Each grey line plots a firm-specific estimated impulse response function. There is a total of 746 firms of which 224 are classified as belonging to group 1 and 522 are classified as belonging to group 2.

Table 1: Frequency of selecting $K = 1, \dots, 5$ when $K^0 = 2$

		Moving Average estimation				
N	T	1	2	3	4	5
100	40	0.01	0.99	0	0	0
100	80	0	1	0	0	0
200	40	0	1	0	0	0
200	80	0	1	0	0	0
		Local Projection estimation				
N	T	1	2	3	4	5
100	40	0.59	0.3	0.11	0	0
100	80	0.3	0.67	0.03	0	0
200	40	0.41	0.29	0.29	0.01	0
200	80	0.1	0.82	0.08	0	0

For each DGP identified by a combination of N and T in the first two columns, this table reports $\frac{1}{250}\sum_{m=1}^{250} \mathbb{1}\{\widehat{K}_m = k\}$ where $\widehat{K}_m$ is the number of groups that minimizes the information criterion defined in (12) for the *m-th* Monte Carlo sample. In the top panel the C-Lasso estimates are obtained from minimising (11) whereas in the bottom panel they are obtained from minimising (15).

Table 2: Comparison of alternative $\beta_{i,4}$ estimators based on Moving Average representation

Group 1										
DGP		Full Heterogeneity			Post Lasso			Group Oracle		
N	T	Bias	Var	MSE	Bias	Var	MSE	Bias	Var	MSE
100	40	0.0015	0.0359	0.0361	-0.0084	0.0058	0.0059	0.0015	0.0008	0.0008
100	80	0.0006	0.0152	0.0153	-0.0018	0.0015	0.0015	0.0006	0.0003	0.0003
200	40	-0.0009	0.0379	0.038	-0.0104	0.0055	0.0056	-0.0009	0.0003	0.0004
200	80	0.0003	0.0151	0.0152	-0.0016	0.0012	0.0012	0.0003	0.0002	0.0002
Group 2										
DGP		Full Heterogeneity			Post Lasso			Group Oracle		
N	T	Bias	Var	MSE	Bias	Var	MSE	Bias	Var	MSE
100	40	0.0001	0.0355	0.0357	0.0099	0.0026	0.0027	0.0001	0.0007	0.0007
100	80	0.0022	0.0153	0.0153	0.0045	0.0005	0.0005	0.0022	0.0004	0.0004
200	40	-0.0006	0.0381	0.0383	0.0089	0.0023	0.0024	-0.0006	0.0004	0.0004
200	80	-0.0003	0.0153	0.0154	0.0016	0.0003	0.0003	-0.0003	0.0001	0.0001

The *Full Heterogeneity* estimator is obtained as the minimizer of (5), the *Post Lasso* estimator is obtained as described in section 4.3 and the *Group Oracle* is obtained as the ex-ante classification estimator from (3) and (4) under the true group membership. For a given estimator $\widehat{\beta}_{i,4}$: (*i*) the *bias* column for group j is computed as $\frac{1}{N_{G_j}} \sum_{i \in G_j} \mathbb{B}_{i,4}$ where $\mathbb{B}_{i,4} = \frac{1}{250} \sum_{m=1}^{250} (\widehat{\beta}_{i,4}^{m} - \beta_{i,4})$ and $\widehat{\beta}_{i,4}^{m}$ denotes the estimates for $\beta_{i,4}$ obtained from the *m-th* Monte Carlo sample of the respective DGP; (*ii*) the *variance* column for group j is computed as $(1/N_{G_j}) \sum_{i \in G_j} \mathbb{V}_{i,4}$ where $\mathbb{V}_{i,4} = (1/250) \sum_{m=1}^{250} (\widehat{\beta}_{i,4}^{m} - \overline{\beta}_{i,4})^2$ where $\overline{\beta}_{i,4} = \frac{1}{250} \sum_{m=1}^{250} \widehat{\beta}_{i,4}^{m}$ and (*iii*) the *MSE* column for group j is computed as $(1/N_{G_j}) \sum_{i \in G_j} \left(\mathbb{B}_{i,4}^2 + \mathbb{V}_{i,4} \right)$.

Table 3: Comparison of alternative $\beta_{i,4}$ estimators based on Local Projections

Group 1										
DGP		Full Heterogeneity			Post Lasso			Group Oracle		
N	T	Bias	Var	MSE	Bias	Var	MSE	Bias	Var	MSE
100	40	0.0003	0.0307	0.0308	-0.0748	0.0094	0.015	0.0003	0.0054	0.0054
100	80	-0.0024	0.0162	0.0162	-0.0335	0.0067	0.0078	-0.0024	0.0034	0.0034
200	40	-0.0001	0.0316	0.0317	-0.0744	0.0096	0.0151	-0.0001	0.0059	0.0059
200	80	0.0131	0.0157	0.0159	-0.0041	0.0057	0.0057	0.0131	0.003	0.0032
Group 2										
DGP		Full Heterogeneity			Post Lasso			Group Oracle		
N	T	Bias	Var	MSE	Bias	Var	MSE	Bias	Var	MSE
100	40	-0.0002	0.0315	0.0316	0.0748	0.0178	0.0234	-0.0002	0.0057	0.0057
100	80	0.0009	0.0161	0.0162	0.0321	0.0101	0.0111	0.0009	0.0033	0.0033
200	40	-0.003	0.0311	0.0312	0.0713	0.0181	0.0232	-0.003	0.0056	0.0056
200	80	-0.0127	0.0155	0.0158	0.0045	0.0077	0.0077	-0.0127	0.0031	0.0033

The *Full Heterogeneity* estimator is obtained as the minimizer of (5), the *Post Lasso* estimator is obtained as described in section 4.3 and the *Group Oracle* is obtained as the ex-ante classification estimator from (3) and (4) under the true group membership. For a given estimator $\widehat{\beta}_{i,4}$: (*i*) the *bias* column for group j is computed as $\frac{1}{N_{G_j}} \sum_{i \in G_j} \mathbb{B}_{i,4}$ where $\mathbb{B}_{i,4} = \frac{1}{250} \sum_{m=1}^{250} (\widehat{\beta}_{i,4}^{m} - \beta_{i,4})$ and $\widehat{\beta}_{i,4}^{m}$ denotes the estimates for $\beta_{i,4}$ obtained from the *m-th* Monte Carlo sample of the respective DGP; (*ii*) the *variance* column for group j is computed as $(1/N_{G_j}) \sum_{i \in G_j} \mathbb{V}_{i,4}$ where $\mathbb{V}_{i,4} = (1/250) \sum_{m=1}^{250} (\widehat{\beta}_{i,4}^{m} - \overline{\beta}_{i,4})^2$ where $\overline{\beta}_{i,4} = \frac{1}{250} \sum_{m=1}^{250} \widehat{\beta}_{i,4}^{m}$ and (*iii*) the *MSE* column for group j is computed as $(1/N_{G_j}) \sum_{i \in G_j} \left(\mathbb{B}_{i,4}^2 + \mathbb{V}_{i,4} \right)$.

Table 4: Firm summary statistics for the two IRF groups

Variable	Group 1	Group 2	p-value
Panel A: Share of Collateral and Flow Borrowers			
Collateral Borrowers	17.05%	22.14%	0.15
Earnings Borrowers	82.95%	77.86%	0.15
Panel B: Average share of Intangible Assets			
Variable	Group 1	Group 2	p-value
Intangible Assets/Total Assets	17.44%	16.89%	0.67
Goodwill/Total Assets	13.31%	12.43%	0.41
Panel C: Firm Size (in billions USD)			
Variable	Group 1	Group 2	p-value
Total Assets	6.64	9.85	0.05
Total Sales	1.73	2.01	0.54
Market capitalization	7.86	9.44	0.47
Panel D: Group Composition by GICS Sectors			
Sector	Group 1	Group 2	p-value
Energy	7.14%	6.32%	0.69
Materials	15.63%	8.24%	0.01
Industrials	30.80%	20.88%	0.01
Consumer Discretionary	16.07%	19.92%	0.20
Consumer Staples	3.57%	9.77%	0.00
Health Care	8.48%	6.13%	0.27
Information Technology	10.71%	8.24%	0.30
Communication Services	0.89%	1.72%	0.33
Utilities	5.80%	18.58%	0.00

The columns *Group 1* and *Group 2* contain the average value of each variable computed across firms that are classified as belonging to groups 1 and 2 by the C-Lasso. The p-value column contains the p-value for the null hypothesis that the mean of a variable in group 1 is equal to the mean in group 2 against a two-sided alternative. The GICS sectors Financials and Real Estate were excluded from the table since there is only one firm in the Financial sector and two firms for real estate in the final sample.

Table 5: Average marginal effects on group membership

	(1)	(2)	(3)	(4)	(5)
Flow Borrower	0.069				0.042
	(0.049)				(0.049)
Intangibles share		0.000			0.000
		(0.001)			(0.001)
Total Assets			-0.002*		-0.002
			(0.001)		(0.002)
Materials				0.085	0.050
				(0.073)	(0.085)
Industrials				0.035	-0.004
				(0.065)	(0.077)
Consumer Disc.				-0.084	-0.109
				(0.069)	(0.079)
Consumer Staples				-0.240**	-0.396***
				(0.094)	(0.137)
Health Care				0.023	-0.033
				(0.081)	(0.100)
IT				0.010	-0.047
				(0.076)	(0.091)
Communication				-0.172	-0.076
				(0.164)	(0.181)
Utilities				-0.271***	-0.212**
				(0.080)	(0.093)
N	578	745	746	746	577

Each column reports the estimated average marginal effects for a Logit specification where the dependent variable is a dummy variable equal to one if the firm belongs to group 1 and 0 if if belongs to group 2. Standard errors in parenthesis. *, ** and *** denote marginal effects that are significant at 10%, 5% and 1% significance levels, respectively.

Appendix to Disaggregated impulse responses via the classifier-Lasso

Miguel Bandeira

A. Proofs of results in the main text

Proof of Proposition 1. Consider first the fully heterogeneous estimator,

$$\widehat{\boldsymbol{\beta}}_i = \left(\mathbf{X}_i'\mathbf{X}_i\right)^{-1}\mathbf{X}_i'\mathbf{y}_i \tag{18}$$

where $\mathbf{y}_i \equiv [y_{i,1}, \ldots, y_{i,T}]'$. The proof of (6) is a standard textbook proof of OLS unbiasedness under the Gauss-Markov assumptions and it follows that $\mathbb{V}\text{ar}\left(\widehat{\boldsymbol{\beta}}_i \mid \mathbf{X}\right) = \sigma^2\left(\mathbf{X}_i'\mathbf{X}_i\right)^{-1}$.[20] The estimator based on the ex-ante group classification approach defined in (3) and (4) is given by,

$$\widetilde{\boldsymbol{\beta}}_i(\tilde{\mathcal{G}}^K) = \left(\sum_{i\in\tilde{G}_a}\mathbf{X}_i'\mathbf{X}_i\right)^{-1}\sum_{i\in\tilde{G}_a}\mathbf{X}_i'\mathbf{y}_i \tag{19}$$

To derive (7) let $\boldsymbol{\varepsilon}_i \equiv [\varepsilon_{i,1}, \ldots, \varepsilon_{i,T}]'$ and use assumptions 1 and 2 to obtain,

[20] See, for instance, Hayashi (2000, section 1.3).

$$
\begin{aligned}
\widetilde{\boldsymbol{\beta}}_i(\tilde{\mathcal{G}}^K) &= \left(\sum_{i\in\tilde{G}_a} \mathbf{X}_i'\mathbf{X}_i\right)^{-1}\left(\sum_{i\in\tilde{G}_a} \mathbf{X}_i'\left(\mathbf{X}_i\boldsymbol{\beta}_i + \boldsymbol{\varepsilon}_i\right)\right) \\
&= \left(\sum_{i\in\tilde{G}_a} \mathbf{X}_i'\mathbf{X}_i\right)^{-1}\left(\sum_{i\in\tilde{G}_a} \mathbf{X}_i'\mathbf{X}_i\boldsymbol{\beta}_i + \sum_{i\in\tilde{G}_a} \mathbf{X}_i'\boldsymbol{\varepsilon}_i\right) \\
&= \left(\sum_{i\in\tilde{G}_a} \mathbf{X}_i'\mathbf{X}_i\right)^{-1}\left(\sum_{i\in\tilde{G}_a}\sum_{k=1}^{K_0} \mathbf{X}_i'\mathbf{X}_i\boldsymbol{\alpha}_k\mathbb{1}\{i\in G_k\} + \sum_{i\in\tilde{G}_a} \mathbf{X}_i'\boldsymbol{\varepsilon}_i\right) \\
&= \left(\sum_{i\in\tilde{G}_a} \mathbf{X}_i'\mathbf{X}_i\right)^{-1}\left(\sum_{i\in\tilde{G}_a\cap G_b} \mathbf{X}_i'\mathbf{X}_i\boldsymbol{\beta}_i + \sum_{\substack{k=1\\k\neq b}}^{K_0}\sum_{i\in\tilde{G}_a\cap G_k} \mathbf{X}_i'\mathbf{X}_i\boldsymbol{\alpha}_k + \sum_{i\in\tilde{G}_a} \mathbf{X}_i'\boldsymbol{\varepsilon}_i\right) \quad (20)
\end{aligned}
$$

Define $\tilde{\boldsymbol{\varphi}}_{a,b} \equiv \left(\sum_{i\in\tilde{G}_a} \mathbf{X}_i'\mathbf{X}_i\right)^{-1}\left(\sum_{i\in\tilde{G}_a\cap G_b} \mathbf{X}_i'\mathbf{X}_i\right)$ and simplify (20) to obtain,

$$
\widetilde{\boldsymbol{\beta}}_i(\tilde{\mathcal{G}}^K) = \tilde{\boldsymbol{\varphi}}_{a,b}\,\boldsymbol{\beta}_i + \sum_{\substack{k=1\\k\neq b}}^{K_0} \tilde{\boldsymbol{\varphi}}_{a,k}\boldsymbol{\alpha}_k + \left(\sum_{i\in\tilde{G}_a} \mathbf{X}_i'\mathbf{X}_i\right)^{-1}\sum_{i\in\tilde{G}_a} \mathbf{X}_i'\boldsymbol{\varepsilon}_i \quad (21)
$$

Taking conditional expectations yields,

$$
\begin{aligned}
\mathbb{E}\left(\widetilde{\boldsymbol{\beta}}_i(\tilde{\mathcal{G}}^K) \mid \mathbf{X}\right) &= \tilde{\boldsymbol{\varphi}}_{a,b}\,\boldsymbol{\beta}_i + \sum_{\substack{k=1\\k\neq b}}^{K_0} \tilde{\boldsymbol{\varphi}}_{a,k}\boldsymbol{\alpha}_k + \left(\sum_{i\in\tilde{G}_a} \mathbf{X}_i'\mathbf{X}_i\right)^{-1}\sum_{i\in\tilde{G}_a} \mathbf{X}_i'\,\mathbb{E}\left(\boldsymbol{\varepsilon}_i \mid \mathbf{X}\right) \\
&= \tilde{\boldsymbol{\varphi}}_{a,b}\,\boldsymbol{\beta}_i + \sum_{\substack{k=1\\k\neq b}}^{K_0} \tilde{\boldsymbol{\varphi}}_{a,k}\boldsymbol{\alpha}_k \quad (22)
\end{aligned}
$$

where the second equality follows from the strict exogeneity in assumption 1. Finally, using the law of iterated expectations we obtain (7),

$$\mathbb{E}\left(\widetilde{\boldsymbol{\beta}}_i(\tilde{\mathcal{G}}^K)\right) = \mathbb{E}\left(\mathbb{E}\left(\widetilde{\boldsymbol{\beta}}_i(\tilde{\mathcal{G}}^K) \mid \mathbf{X}\right)\right) = \boldsymbol{\varphi}_{a,b}\,\boldsymbol{\beta}_i + \sum_{\substack{k=1\\k\neq b}}^{K_0} \boldsymbol{\varphi}_{a,k}\boldsymbol{\alpha}_k \tag{23}$$

where $\boldsymbol{\varphi}_{a,b} \equiv \mathbb{E}\left(\tilde{\boldsymbol{\varphi}}_{a,b}\right)$. Proving (8) requires showing that $\mathbb{V}\text{ar}\left(\widehat{\boldsymbol{\beta}}_i \mid \mathbf{X}\right) - \mathbb{V}\text{ar}\left(\widetilde{\boldsymbol{\beta}}_i(\tilde{\mathcal{G}}^K) \mid \mathbf{X}\right)$ is positive semi-definite. Start by using (20) to derive $\mathbb{V}\text{ar}\left(\widetilde{\boldsymbol{\beta}}_i(\tilde{\mathcal{G}}^K) \mid \mathbf{X}\right)$,

$$\begin{aligned}
\mathbb{V}\text{ar}\left(\widetilde{\boldsymbol{\beta}}_i(\tilde{\mathcal{G}}^K) \mid \mathbf{X}\right) &= \left(\sum_{i\in\tilde{G}_a} \mathbf{X}_i'\mathbf{X}_i\right)^{-1} \mathbb{V}\text{ar}\left(\sum_{i\in\tilde{G}_a} \mathbf{X}_i'\varepsilon_i\right)\left(\sum_{i\in\tilde{G}_a} \mathbf{X}_i'\mathbf{X}_i\right)^{-1} \\
&= \left(\sum_{i\in\tilde{G}_a} \mathbf{X}_i'\mathbf{X}_i\right)^{-1} \left(\sum_{i\in\tilde{G}_a} \mathbf{X}_i'\mathbb{V}\text{ar}\left(\varepsilon_i \mid \mathbf{X}\right)\mathbf{X}_i\right)\left(\sum_{i\in\tilde{G}_a} \mathbf{X}_i'\mathbf{X}_i\right)^{-1} \\
&= \sigma^2\left(\sum_{i\in\tilde{G}_a} \mathbf{X}_i'\mathbf{X}_i\right)^{-1}
\end{aligned} \tag{24}$$

where the second and third equalities follow from the conditional homoskedasticity and no autocorrelation assumption. Combining the expressions for $\mathbb{V}\text{ar}\left(\widehat{\boldsymbol{\beta}}_i \mid \mathbf{X}\right)$ and $\mathbb{V}\text{ar}\left(\widetilde{\boldsymbol{\beta}}_i(\tilde{\mathcal{G}}^K) \mid \mathbf{X}\right)$ yields,

$$\mathbb{V}\text{ar}\left(\widehat{\boldsymbol{\beta}}_i \mid \mathbf{X}\right) - \mathbb{V}\text{ar}\left(\widetilde{\boldsymbol{\beta}}_i(\tilde{\mathcal{G}}^K) \mid \mathbf{X}\right) = \sigma^2 \underbrace{\left(\left(\mathbf{X}_i'\mathbf{X}_i\right)^{-1} - \left(\textstyle\sum_{i\in\tilde{G}_a} \mathbf{X}_i'\mathbf{X}_i\right)^{-1}\right)}_{[*]} \tag{25}$$

The term $[*]$ is positive semi definite if and only if $\sum_{i\in\tilde{G}_a} \mathbf{X}_i'\mathbf{X}_i - \mathbf{X}_i'\mathbf{X}_i$ is positive semi definite. Finally,

$$\textstyle\sum_{i\in\tilde{G}_a} \mathbf{X}_i'\mathbf{X}_i - \mathbf{X}_i'\mathbf{X}_i = \sum_{j\in\tilde{G}_a\setminus\{i\}} \mathbf{X}_j'\mathbf{X}_j \tag{26}$$

If $\tilde{G}_a \backslash \{i\} = \varnothing$ then $\sum_{i \in \tilde{G}_a} \mathbf{X}_i' \mathbf{X}_i - \mathbf{X}_i' \mathbf{X}_i = \mathbf{0}$. If $\tilde{G}_a \backslash \{i\} \neq \varnothing$, then $\sum_{i \in \tilde{G}_a} \mathbf{X}_i' \mathbf{X}_i - \mathbf{X}_i' \mathbf{X}_i$ is the sum of positive definite matrices and, hence, positive definite. Therefore, $[*]$ is positive semi-definite. $\square$