

Giacomini, Raffaella; Lee, Sokbae; Sarpietro, Silvia

Working Paper

A robust method for microforecasting and estimation of random effects

Working Paper, No. WP 2023-26

Provided in Cooperation with:

Federal Reserve Bank of Chicago

Suggested Citation: Giacomini, Raffaella; Lee, Sokbae; Sarpietro, Silvia (2023) : A robust method for microforecasting and estimation of random effects, Working Paper, No. WP 2023-26, Federal Reserve Bank of Chicago, Chicago, IL, <https://doi.org/10.21033/wp-2023-26>

This Version is available at:

<https://hdl.handle.net/10419/284067>

Standard-Nutzungsbedingungen:

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

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

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

Terms of use:

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

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

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

A Robust Method for Microforecasting and Estimation of Random Effects

Raffaella Giacomini, Sokbae Lee,
and Silvia Sarpietro

August 2, 2023

WP 2023-26

<https://doi.org/10.21033/wp-2023-26>

FEDERAL RESERVE BANK *of* CHICAGO

*Working papers are not edited, and all opinions are the responsibility of the author(s). The views expressed do not necessarily reflect the views of the Federal Reserve Bank of Chicago or the Federal Reserve System.

A Robust Method for Microforecasting and Estimation of Random Effects^{*}

Raffaella GIACOMINI [†] Sokbae LEE [‡] Silvia SARPIETRO [§]

August 2, 2023

Abstract

We propose a method for forecasting individual outcomes and estimating random effects in linear panel data models and value-added models when the panel has a short time dimension. The method is robust, trivial to implement and requires minimal assumptions. The idea is to take a weighted average of time series- and pooled forecasts/estimators, with individual weights that are based on time series information. We show the forecast optimality of individual weights, both in terms of minimax-regret and of mean squared forecast error. We then provide feasible weights that ensure good performance under weaker assumptions than those required by existing approaches. Unlike existing shrinkage methods, our approach borrows the strength - but avoids the tyranny - of the majority, by targeting individual (instead of group) accuracy and letting the data decide how much strength each individual should borrow. Unlike existing empirical Bayesian methods, our frequentist approach requires no distributional assumptions, and, in fact, it is particularly advantageous in the presence of features such as heavy tails that would make a fully nonparametric procedure problematic.

KEYWORDS: Minimax-Regret; Shrinkage; Forecast Combination; Robustness

JEL CLASSIFICATION: C10, C23, C53

^{*}We are grateful to Roger Koenker, Laura Liu, Mikkel Plagborg-Møller and Suyong Song for valuable comments and suggestions. We benefited from useful comments by seminar participants at the 28th International Panel Data Conference (Amsterdam), Dolomiti Macro Meeting, US Panel Data Conference, CSEF-DISES (Naples), BI Norwegian Business School, IWEER 2022 (Rimini), Duke Microeconometrics Class of 2020 and 2021 Conference, ICEEE 2023 (Cagliari). The views expressed here are those of the authors and do not reflect the views of the Federal Reserve Bank of Chicago or the Federal Reserve System.

[†]Chicago Fed and University College London. r.giacomini@ucl.ac.uk

[‡]Columbia University. sl3841@columbia.edu

[§]University of Bologna. silvia.sarpietro@unibo.it

1 Introduction

*“Knowing when to borrow and when not to borrow is one of the key aspects of statistical practice”
(Mallows and Tukey [1982])*

Estimation of individual effects (i.e., fixed or random effects) and microforecasting (i.e., forecasting individual outcomes using panel data with a short time dimension) are of prominent interest in empirical economics. For example, Kline, Rose, and Walters [2022] estimate the distribution of firm-specific effects in order to analyze bias in firms’ hiring decisions. Estimating individual effects and forecasting are the goals of the literature using “value added models” to capture institutional effects, e.g., the effect of teachers on students’ test scores [Kane and Staiger, 2008, Chetty, Friedman, and Rockoff, 2014a,b, Angrist, Hull, Pathak, and Walters, 2017], the effect of neighborhoods on intergenerational mobility [Chetty and Hendren, 2018], or the effect of physicians on patients’ outcomes [Fletcher, Horwitz, and Bradley, 2014]. As discussed by Hull [2020], these estimates and forecasts play an important policy role in the regulation of healthcare and education in the U.S. Other examples in microeconomics include the literature on long-term treatment effects, which relies on forecasting the effects of treatments such as early-childhood interventions [García, Heckman, Leaf, and Prados, 2020] or job-training programs [Athey, Chetty, Imbens, and Kang, 2019], and the literature that forecasts individual incomes for consumption/savings decisions [Chamberlain and Hirano, 1999]. Macroeconomic panel forecasting also falls into this category if it uses short estimation windows to account for parameter instability (e.g. Liu, Moon, and Schorfheide [2020] forecast banks’ revenues after a regulatory change).

The short time dimension in the panel introduces a bias-variance tradeoff: estimators/forecasts based on individual time series capture individual effects but are noisy; pooling information reduces the variance but results in bias. A leading approach in the literature (e.g., the value-added literature discussed above, Kline et al. [2022], and Liu et al. [2020]) is to rely on Bayesian methods that “borrow strength” from the majority in order to improve accuracy (Mallows and Tukey [1982]). These approaches however present two main drawbacks: they suffer from the “tyranny of the majority” or they impose distributional assumptions.

An example of the former are shrinkage estimators, such as the classical James and Stein [1961]’s estimator (which assumes homogeneous variances of the random effects and of the idiosyncratic shocks), or the extension considered by Kwon [2021] (which assumes homogeneous variance for the random effects but allows for heterogeneous variance of the idiosyncratic shocks).¹ Shrinkage estimators are typically adopted in the value-added literature discussed above. Shrinkage estimators are known to suffer from what Efron [2010] calls the “tyranny of the majority”. Intuitively, they shrink individuals by the same amount regardless of their ran-

¹Pesaran, Pick, and Timmermann [2022] consider a shrinkage approach in a linear model with heterogeneous parameters on covariates and similarly impose a homogeneous variance assumption for the parameters.

dom effects, which penalizes individuals with random effects that are far away from the common mean - outliers - as well as individuals with random effects near the common mean, as we show in Section 6).² At the core of the phenomenon lies the fact that existing shrinkage approaches target group accuracy rather than individual accuracy. Since individual accuracy is typically the object of interest in the empirical literature, we argue that it is crucial to consider methods that deliver accuracy at the individual level first and foremost.³ An insight of this paper is that overcoming the tyranny of the majority can be linked to relaxing the parameter homogeneity assumption made by existing shrinkage approaches, which yields weights that implicitly relate the amount of shrinkage to how far the random effect is from the common mean. Allowing for fully heterogeneous parameters is thus a crucial ingredient for obtaining accuracy at the individual level, as well as being empirically relevant (see, e.g., the discussion in Kwon [2021] for the motivation behind allowing for heterogeneous variance of the idiosyncratic shocks).

A recent literature adopts a different approach to microforecasting and estimation of random effect, based on empirical Bayes estimators. This approach is considered by, e.g., Gu and Koenker [2017] and Liu et al. [2020] or in the application by Kline et al. [2022] using the empirical Bayes deconvolution estimator proposed by Efron [2010]. These estimators rely on a parametric distributional assumption for the idiosyncratic shocks. This can be considered restrictive and can make it challenging to incorporate features such as distributions with heavy tails, which are empirically relevant, e.g., in the case of wealth or earnings outcomes (Güvenen, Karahan, Ozkan, and Song [2015], Browning, Ejrnaes, and Alvarez [2010]). In principle, it is possible to relax the parametric assumption by considering nonparametric deconvolution estimators (e.g. Delaigle, Hall, and Meister [2008]), however the convergence rate of these estimators depends on tail behavior and tends to be slow (see, e.g., Fan [1991] and Fan and Truong [1993]). In addition, the practical implementation of these estimators requires integration and regularization due to the ill-posed inverse problem and is known to present considerable difficulties (see, e.g., Delaigle and Gijbels [2007], Delaigle [2014] and Hall and Meister [2007]).⁴

In this paper, we instead go back to the classical approach and focus our attention on shrinkage methods, which do not rely on distributional assumptions and are trivial to implement. We present a method that overcomes the tyranny of the majority problem that affects existing shrinkage approaches. The method also turns heavy tails from vice to virtue, in the sense that it is particularly advantageous relative to existing methods in the presence of heavy tails.

We consider a frequentist, finite-sample setting under minimal assumptions and propose an

²Attempts to tackle the problem by first identifying and then not shrinking outliers (Efron and Morris [1971]) are complicated by the need to specify an arbitrary threshold and by the presence of different factors affecting individual performance. For example, it is not just outliers in terms of the random effects who benefit less from borrowing strength, but also “creatures of habit” - individuals with low variance of the idiosyncratic shocks for whom the time-series forecast/estimator is therefore not very noisy.

³The value-added literature, for example, is often interested in identifying outliers and performing policy counterfactuals that involve them.

⁴Regularization is also recommended by Efron [2010] to ensure accuracy, which introduces another element of choice and a source of sensitivity even when the parametric assumption is correct.

approach to microforecasting and estimation of random effects that targets individual accuracy. The general idea is to consider a weighted average of estimators/forecasts in a given class, using individual-specific weights that exploit the information contained in the time series dimension (even if it is very short). We call our approach Individual Weighting (henceforth IW). When the class includes time series- (TS) and pooled (Pool) estimators/forecasts, IW gives a shrinkage estimator that differs from existing shrinkage approaches by leveraging the time series- instead of the cross-sectional dimension to compute weights. This is what allows IW to borrow the strength - but avoid the tyranny - of the majority, by letting past time-series data decide how much strength each individual should borrow. The downside of leveraging only the small time series dimension, instead of the potentially large cross-sectional dimension, is that we cannot rely on asymptotic behaviour to judge the performance of our method and to construct standard errors. Instead, we focus on finite-sample optimality of forecast performance and on robustness, that is, on obtaining a method that performs well over the unknown parameter space. Our approach is inspired by Manski [2021], who emphasizes evaluation of decision rules by their performance across the parameter space and advocates the minimax regret criterion.⁵

We focus on a model with no distributional assumptions, where individual outcomes are the sum of random effects and idiosyncratic shocks, both with heterogeneous variances. The random effects have a common mean, which is a tuning parameter representing how we borrow strength from the majority. It is assumed to be either known (if, for example, the outcomes are demeaned) or it is approximated with the pooled mean over the panel. Outcomes can be also redefined as residuals from the first-step estimation of a linear panel data model or a value-added model with homogeneous coefficients for the covariates (possibly including lagged outcomes). Our setting thus encompasses a large class of empirically relevant models. Random effects and shocks are assumed to be independent of each other, however we do not restrict the relationship between the parameters characterizing their distribution.

The theoretical results focus on forecasting and show conditions under which IW with generic weights is optimal relative to using TS or Pool for all individuals. No forecast uniformly dominates the others over the parameter space in terms of individual accuracy, as measured by the Mean Squared Forecast Error (MSFE). The advantage of IW is however two-fold: it is robust - it avoids large errors in regions of the parameter space where the accuracies of TS and Pool are different - and it delivers improvements when TS and Pool are equally accurate.

Formally, we show that IW is minimax-regret optimal over the parameter space relative to using either TS or Pool. In addition, IW is also MSFE optimal if we restrict attention to the region of the parameter space where TS and Pool are equally accurate. The improvement of IW in terms of both minimax-regret and MSFE is strict for any constant weight between 0 and 1. Furthermore, keeping all else equal, an additional improvement can be obtained when the

⁵There are a couple of papers that apply the minimax regret criterion to panel data: handling missing data in sample design [Dominitz and Manski, 2021] and forecasting discrete outcomes under partial identification or other concerns [Christensen, Moon, and Schorfheide, 2020]. Their focus is distinct from ours.

weights satisfy a key assumption requiring that: 1) the weights are genuine functions of the random effect; and 2) the weights are not “pathological”, in the sense that they do not shrink outliers more than random effects near the mean of the distribution (which would exacerbate the tyranny of the majority phenomenon). Existing shrinkage approaches, for example, deliver weights that are strictly between 0 and 1 but that are based on cross-sectional information and thus are not functions of the random effects. This means that shrinkage approaches outperform using TS or Pool for all individuals, but can be outperformed by weights that satisfy the key assumption. Finally, we show that the magnitude of the improvement over TS and Pool for weights that satisfy the key assumption is larger the heavier the tails of the distribution of random effects. In other words, heavy tails confer an advantage to shrinkage approaches that are able to capture the random effect via individual-specific weights.

The robustness of our approach is rooted in the focus on individual- rather than group accuracy. It is however worth noting that IW can end up dominating existing methods also in terms of group accuracy. Whether this occurs in a given application depends on the distribution of the parameters across individuals - which our individual-level analysis purposely leaves unrestricted - and also on the unknown tail properties of the distribution of the random effects.

We present three types of feasible weights that satisfy the key assumption discussed above. The first type of feasible weights are based on oracle weights that are optimal in terms of the individual MSFE. These oracle weights correspond to the classical Bates and Granger [1969]’s weights applied here to the individual MSFE and are a function of the heterogeneous variance parameters. Feasible weights can then be obtained by estimating these parameters using the time series dimension. These “estimated oracle individual weights” (IW-O) are likely to perform poorly due to the short time dimension. In the search for feasible weights that perform well, and that do so over the parameter space, we further consider a minimax regret analysis where we condition the individual MSFE on the past. This analysis allows us to characterize optimal weights as functions of certain conditional expectations. We then show that bounds on a particular conditional expectation translate into optimal weights that have a feasible counterpart. We call this second type of feasible weights “minimax-regret optimal individual weights” (IW-MR). Finally, we consider alternative feasible weights based on the individual inverse squared forecast error (IW-MSFE), which are equivalent to the weights considered in some forecast combination time-series literature (e.g., Stock and Watson [1998]) but computed on a very short time series. These weights offer additional robustness benefits because they do not rely on correct specification of the model and can thus be applied in more general settings.

We compare the finite sample performance of all the feasible weights that we consider in this paper, and conclude that IW-MR are the preferred weights. Additional simulations compare IW to Bayesian shrinkage estimators. We illustrate how IW can overcome the tyranny of the majority phenomenon that affects existing shrinkage approaches (specifically, the James-Stein estimator): IW tends to deliver accuracy gains for individuals with random effects in the tails

and near the mean of the distribution, the more so the heavier the tails and the larger the variance of the distribution. This implies that IW can also outperform James-Stein in terms of group accuracy, again depending on the tail properties of the random effects distribution, in relation to its variance.

We consider two empirical illustrations. The first analyzes gender discrimination in firm hiring as in Kline et al. [2022]. IW yields an estimated distribution of random effects that is slightly more right-skewed and has a heavier right tail than the distribution obtained in the paper (which is based on Efron [2010]), implying a similar measure of overall inequality but a larger proportion of firms discriminating in favour of male applicants. IW also delivers more accurate out-of-sample forecasts for firms near the center and in the right tail of the distribution, and also in terms of group accuracy. The second application forecasts individual earning residuals out-of-sample using the Panel Study of Income Dynamics. We find that the forecast with the best group accuracy is IW, which tends to assign high weights to pooling for individuals with earning residuals near the median of the distribution. This application illustrates the usefulness of IW even in highly heterogeneous settings where only a few individuals really benefit from borrowing strength from the majority.

The rest of the paper is organized as follows. Section 2 introduces IW in a general setting. Section 3 considers a model with random effects and shows how to implement IW in practice. Section 4 shows the optimality of IW, both in terms of minimax-regret and in terms of Mean Squared Forecast Error. Section 5 derives the three different types of feasible weights that we propose in the paper. Section 6 shows simulation results. Section 7 illustrates the performance of IW in two empirical applications. Section 8 offers concluding remarks. Appendix A contains the proofs and Appendix B derives the feasible weights under more general assumptions than those considered in section 5 for illustrative purposes.

2 IW in General

We start by outlining the general problem. For each individual i , we have a class of K different forecasts/estimators:

$$\mathcal{F} = \{\hat{Y}_{k,i,T}, k = 1, \dots, K\}.$$

Each forecast/estimator is a function of the information set at time T , $\mathcal{Y}_{N,T} := \{Y_{i,t} : i = 1, \dots, N, t = 1, \dots, T\}$.⁶ The panel has a short T .

For each i , IW considers a weighted average of the forecasts/estimators in the class with

⁶For simplicity of notation, we focus on a balanced panel, but the extension to unbalanced panels is straightforward.

individual weights $W_{k,i,T}$ that depend on $\mathcal{Y}_{N,T}$:

$$\hat{Y}_{i,T}^{IW} = \sum_{k=1}^K \hat{Y}_{k,i,T} W_{k,i,T}. \quad (1)$$

A byproduct of the procedure that may be of interest in some applications is a fuzzy clustering of individuals, grouped according to which forecast/estimator receives the largest weight.

3 IW in a Model with Random Effects

In the remainder of the paper, we specialize the discussion to a model with random effects where IW combines a time series forecast/estimator and a pooled forecast/estimator. Since in this model the estimator of the random effect coincides with the forecast of the individual outcome, for simplicity we focus the discussion in the rest of the paper on forecasting.

3.1 The Model

The baseline model assumes that the outcomes are the sum of individual random effects and idiosyncratic shocks, both with heterogeneous variances:

$$Y_{i,t} = A_i + U_{i,t}, i = 1, \dots, N; t = 1, \dots, T, \quad (2)$$

where $A_i \sim (\mu, \lambda_i^2)$ and $U_{i,t} \sim (0, \sigma_i^2)$. Here, $A_i, U_{i,1}, \dots, U_{i,T}$ are random variables, whereas μ, λ_i^2 and σ_i^2 are parameters. In other words, we take the frequentist approach.

We make the following assumption.

Assumption 3.1 (Independence). $A_i, U_{i,1}, \dots, U_{i,T}$ are mutually independent.

Remark 1 (Relationship of assumptions with Bayesian approaches). As a point of comparison, we will relate our approach to two classes of Bayesian estimators: empirical Bayes deconvolution estimators (in our simulations and empirical application we will focus on Efron [2010]) and shrinkage estimators such as James and Stein [1961] and Kwon [2021]. In the context of the model in (2), Efron [2010] amounts to imposing a normality assumption on the idiosyncratic shocks, $U_{i,t} \sim N(0, \sigma^2)$, whereas $A_i \sim G_A(\mu, \lambda^2)$, with G_A flexibly parameterized by a spline. James and Stein [1961] and Kwon [2021] do not require distributional assumptions but impose homogeneity assumptions on the variance parameters, respectively assuming $U_{i,t} \sim (0, \sigma^2)$, $A_i \sim (\mu, \lambda^2)$ and $U_{i,t} \sim (0, \sigma_i^2)$, $A_i \sim (\mu, \lambda^2)$. Like existing approaches, we assume a common μ across individuals, but we differ by allowing all other parameters to be individual-specific and by making no distributional assumptions.

Remark 2 (Interpretation of μ). The parameter μ represents how we “borrow strength” from the majority. In practice, we do this by shrinking the time-series forecasts/estimators of the random effects A_i towards their mean μ . μ plays a similar role in our analysis as in a classical Bayesian setting, and we similarly consider it as a tuning parameter. As discussed by Kwon [2021], in empirical work the outcomes are often first demeaned, in which case one simply sets $\mu = 0$ (for example if $Y_{i,t}$ are residuals from a first-stage estimation of a model with an intercept, see remark 3 below). If μ is unknown, we replace it with the sample mean of $Y_{i,t}$ over the panel. See remark 8 below for a discussion of how μ could be chosen in case of a known group structure in parameters. The assumption of a common μ across individuals is thus only needed when the point of shrinkage is not known, so that one can leverage the panel dimension to estimate it. Treating the point of shrinkage as a tuning parameter means that, similarly to existing approaches, our theoretical results do not take into account possible uncertainty in its estimation.

Remark 3 (Extension: covariates). Covariates $X_{i,t}$ can be incorporated by redefining $Y_{i,t}$ in (2) as residuals from the first-step estimation of a model with homogeneous coefficients:

$$Y_{i,t} = \tilde{Y}_{i,t} - X'_{i,t}\hat{\beta}, \quad (3)$$

where $\tilde{Y}_{i,t}$ are the outcomes and $\hat{\beta}$ is a consistent estimator of the homogeneous coefficients as $N \rightarrow \infty$.⁷ All the theoretical results discussed below then apply under the additional assumption that N is large. Note that the assumption of consistency of $\hat{\beta}$ could in principle be relaxed, as in a finite- N setting there can be other, perhaps biased, estimators that improve forecast accuracy. We leave this extension for future research. The extension to a model with heterogeneous coefficients for the covariates would imply generalizing the problem of estimating/forecasting unobserved heterogeneity from the univariate case considered here to the multivariate case. While a full treatment of this extension is beyond the scope of this paper, we offer some remarks on this topic in the conclusion of the paper.

Remark 4 (Extension: value-added model). A value added model for an outcome $\tilde{Y}_{i,j,t}$ and covariates $X_{i,j,t}$ aims at estimating the random effect A_i in the model

$$\tilde{Y}_{i,j,t} = X'_{i,j,t}\beta + A_i + U_{i,j,t}.$$

In the case of teacher value-added, for example, i is the teacher and $j = 1, \dots, n_{i,t}$ (with $n_{i,t}$ finite) are the students assigned to teacher i in year t . This model can also be nested in (2) if a consistent estimator (as $N \rightarrow \infty$) $\hat{\beta}$ is available, in which case one defines $Y_{i,t}$ as residuals from

⁷For instance, if the covariates $X_{i,t}$ include the lagged outcome, one could use the Arellano-Bond estimator [Arellano and Bond, 1991] to estimate β .

a first step estimation involving averaged outcomes and covariates, e.g.,

$$Y_{i,t} = \frac{1}{n_{i,t}} \sum_{j=1}^{n_{i,t}} \tilde{Y}_{i,j,t} - \frac{1}{n_{i,t}} \sum_{j=1}^{n_{i,t}} X'_{i,j,t} \hat{\beta}. \quad (4)$$

Remark 5 (Robustness to distributional assumptions). Note that we make no distributional assumptions on the random effects and the idiosyncratic shocks. Heavy tails in both distributions are permitted, as long as the variances exist (i.e., the parameters λ_i^2 and σ_i^2 are finite).

Remark 6 (Robustness to dependence structure). The analysis below is carried out at the individual level, and thus in principle does not require a large N nor any restriction on the cross-sectional dependence of $Y_{i,t}$. This is true as long as the point of shrinkage μ is known (e.g., when the outcomes have been demeaned so that $\mu = 0$). When the group mean μ is unknown, in practice we will approximate it with the sample mean over the panel, which implicitly requires a restriction on the cross-sectional dependence that ensures validity of a law of large numbers. The incorporation of covariates discussed in remark 3 also implicitly restricts the dependence structure by assuming availability of a consistent estimator of the homogeneous coefficients. Time-series dependence can be accounted for by including lagged dependent variables as covariates, as long as the autoregressive coefficients are homogeneous across individuals.

Remark 7 (Robustness to distribution of parameters across individuals). Our analysis is carried out at the individual level and is purposely agnostic about the distribution of λ_i^2 and σ_i^2 across i . This implies that, in general, we cannot make any formal statement about the group accuracy of our estimator. Nonetheless, we are able to provide some intuition for the implications of our findings for group accuracy, see Section 4.6 below. Another implication is that, while we assume independence between individual random effects and shocks, we accommodate any type of unknown relationship between their variances (e.g., there could be two groups of individuals, one with low λ_i^2 and low (high) σ_i^2 and one with high λ_i^2 and high (low) σ_i^2). The next remark highlights how the analysis can be modified if one is willing to assume a known group structure in parameters.

Remark 8 (Known group structure in parameters). Our analysis assumes fully heterogeneous parameters λ_i^2 and σ_i^2 , and shrinks the estimators/forecasts of the random effects towards the mean μ of the random effects. Suppose there is a group structure in μ only, with a finite number of subgroups and observable group membership (but with λ_i^2 and σ_i^2 still heterogeneous within the subgroups). In this case, the only modification to our analysis is that the point of shrinkage is the mean for the subgroup instead of the mean for the whole panel. If the homogeneity assumption within subgroups extends to λ_i^2 and σ_i^2 , then our estimator becomes the James-Stein estimator applied to each subgroup (and thus it is exactly the James Stein estimator if there is only one group).

Henceforth, we focus on model (2), with the understanding that $Y_{i,t}$ are either outcomes or residuals such as (3) or (4) (in a large- N setting).

3.2 IW in Practice

In practice, for each individual i , we propose considering a weighted average of the time-series and pooled forecasts at time T , with weights based on time series information. The forecasts are:

$$\text{Time series forecast (TS)} : \hat{Y}_{i,T}^{TS} = \bar{Y}_{i,T} = \Sigma_{t=1}^T Y_{i,t}/T, \quad (5)$$

$$\text{Pooled forecast (Pool)} : \hat{Y}_{i,T}^{Pool} = \mu, \quad (6)$$

$$\text{Weighted forecast (IW)} : \hat{Y}_{i,T}^{IW} = \hat{Y}_{i,T}^{TS} W_{i,T} + \hat{Y}_{i,T}^{Pool} (1 - W_{i,T}), \quad (7)$$

where $W_{i,T}$ is one of the feasible individual weights reported below. The parameter μ is either known (e.g. $\mu = 0$ if the observations have been demeaned or if $Y_{i,t}$ are residuals from a first step estimation that includes an intercept) or it is approximated with the pooled mean, $\mu \approx \Sigma_{i=1}^N \Sigma_{t=1}^T Y_{i,t}/NT$.

We derive three classes of feasible individual weights in a simplified setting in Section 5 and their generalizations in Appendix B. For convenience, we summarize here the three classes of general feasible individual weights that we discuss in the paper.

3.2.1 Estimated Oracle Weights (IW-O)

These weights are based on estimated oracle weights, and are given by

$$W_{i,T}^{IW-O} = \left(\frac{(\Sigma_{t=1}^T (Y_{i,t} - \mu)^2/T - \Sigma_{t=1}^{T-1} (Y_{i,t} - Y_{i,t+1})^2/2(T-1))^+}{\Sigma_{t=1}^T (Y_{i,t} - \mu)^2/T - \Sigma_{t=1}^{T-1} (Y_{i,t} - Y_{i,t+1})^2/2T} \right)^+, \quad (8)$$

where $(\cdot)^+$ denotes the positive part.

3.2.2 Minimax Regret Optimal weights (IW-MR)

These weights are given by:⁸

$$W_{i,T}^{IW-MR} = 1 - \frac{1}{\sqrt{\frac{\max\{(Y_{i,1}-\mu)^2, \dots, (Y_{i,T}-\mu)^2\}}{\sum_{t=1}^{T-1} (Y_{i,t}-Y_{i,t+1})^2 / 2T(T-1)} + 1}}. \quad (9)$$

These are the weights that perform best in our simulations.

3.2.3 Inverse MSFE Weights (IW-MSFE)

These weights are based on the inverse MSFE. Since they do not rely on the model assumptions, these weights are applicable in the general setting considered in section 2. The in-sample inverse MSFE weights are:

$$W_{i,T}^{IW-MSFE-IS} := \frac{1 / \left[\sum_{t=1}^T (Y_{i,t} - \hat{Y}_{i,T}^{TS})^2 \right]}{1 / \left[\sum_{t=1}^T (Y_{i,t} - \hat{Y}_{i,T}^{TS})^2 \right] + 1 / \left[\sum_{t=1}^T (Y_{i,t} - \hat{Y}_{i,T}^{Pool})^2 \right]}. \quad (10)$$

For a given choice of $P < T$, the out-of-sample inverse MSFE weights are:

$$W_{i,T,P}^{IW-MSFE-OOS} := \frac{1 / \left[\sum_{t=T-P+1}^T (Y_{i,t} - \hat{Y}_{i,t-1}^{TS})^2 \right]}{1 / \left[\sum_{t=T-P+1}^T (Y_{i,t} - \hat{Y}_{i,t-1}^{TS})^2 \right] + 1 / \left[\sum_{t=T-P+1}^T (Y_{i,t} - \hat{Y}_{i,t-1}^{Pool})^2 \right]}, \quad (11)$$

where $\hat{Y}_{i,t-1}^{TS}$ is the time series mean using data up to time $t-1$ (using all the data available or just an arbitrary number of most recent data) and $\hat{Y}_{i,t-1}^{Pool} = \mu$ (with μ known or approximated with the panel mean on data up to time $t-1$).

4 Optimality of IW

In this section we show conditions under which IW is minimax regret optimal relative to using the TS or Pool forecast for all individuals. We further show conditions under which IW is also optimal in terms of MSFE.

⁸A similar performance in simulations and in the empirical applications is obtained by the IW-MR rule based on an alternative unbiased estimator for σ_i^2 , as follows:

$$W_{i,T}^{IW-MR2} = 1 - \frac{1}{\sqrt{\frac{\max\{(Y_{i,1}-\mu)^2, \dots, (Y_{i,T}-\mu)^2\}}{\sum_{t=1}^T (Y_{i,t}-Y_{i,T})^2 / T(T-1)} + 1}}.$$

4.1 Set Up and Key Assumption

This section considers a simplified setting where the time series forecast is the time- T outcome and IW is based on data available at time $T - 1$. These assumptions imply that the weights and the forecasts are independent, which makes the theoretical results transparent and intuitive. This assumption does not affect the practical usefulness of the results, as the feasible weights that we recommend to use and that we reported in section 3.2 above consider the general case where TS is the time series mean and the weights are based on the information set at time T . The simulations and empirical applications are also based on the general weights in section 3.2. In this section we thus let:

$$\text{Time series forecast (TS)} : \hat{Y}_{i,T}^{TS} = Y_{i,T}, \quad (12)$$

$$\text{Pooled forecast (Pool)} : \hat{Y}_{i,T}^{Pool} = \mu, \quad (13)$$

$$\text{Weighted forecast (IW)} : \hat{Y}_{i,T}^{IW} = Y_{i,T}W_{i,T-1} + \mu(1 - W_{i,T-1}), \quad (14)$$

where $W_{i,T-1}$ depends on data available at time $T - 1$.

Consider a situation where the user is uncertain about the parameter $\theta_i = (\lambda_i^2, \sigma_i^2)$. The unconditional MSFE of forecast m for a given θ_i is

$$\text{MSFE}(m, \theta_i) = \mathbb{E} \left[\left(Y_{i,T+1} - \hat{Y}_{i,T}^m \right)^2 \right],$$

with $m \in \mathcal{M}$.

The next lemma derives the MSFEs of TS, Pool and IW.

Lemma 4.1. *Consider the forecasts in (12), (13), and (14). Then under Assumption 3.1 we have*

$$\begin{aligned} \text{MSFE}(\text{TS}, \theta_i) &= 2\sigma_i^2, \\ \text{MSFE}(\text{Pool}, \theta_i) &= \lambda_i^2 + \sigma_i^2, \\ \text{MSFE}(\text{IW}, \theta_i) &= \sigma_i^2 + \sigma_i^2 \mathbb{E} [W_{i,T-1}^2] + \mathbb{E} [(A_i - \mu)^2 (1 - W_{i,T-1})^2]. \end{aligned}$$

Lemma 4.1 suggests that the trade-off between TS and Pool depends on the “signal-to-noise” ratio λ_i^2/σ_i^2 : Pool dominates when the ratio is less than 1 and TS dominates when it is greater than 1. Knowledge of the parameters would allow one to choose the best forecast; however in the presence of uncertainty about the parameters it is not possible to choose a forecast optimally. We thus pursue an alternative route. We look for a robust rule that performs well over the entire parameter space, in the sense of avoiding large errors when TS and Pool have different accuracy and improving on both TS and Pool when they have similar accuracy. The following two sections show that IW can accomplish both goals. We first formalize the notion of robustness that we consider here, based on the minimax regret criterion.

Let \mathcal{M} include TS, Pool, and IW. We define regret as

$$R(m, \theta_i) := \text{MSFE}(m, \theta_i) - \min_{h \in \mathcal{M}} \text{MSFE}(h, \theta_i). \quad (15)$$

The minimax-regret (MMR) criterion selects the forecast m that minimizes the maximum regret

$$\max_{\theta_i \in \Theta} R(m, \theta_i),$$

where Θ is the parameter space. The form of regret here is similar to that of regret in decision theory without sample data (e.g., see equation (3) in Manski [2021]). The MMR criterion is championed by Manski [2021].⁹ The regret in (15) is defined relative to the best MSFE out of only three models because the goal in this section is to choose among IW, TS, and Pool.¹⁰

To derive analytical results for IW, we impose the following key regularity condition.

Assumption 4.1 (Individual Weight). *The individual weight $W_{i,T-1}$ satisfies $0 \leq W_{i,T-1} \leq 1$ and*

$$\mathbb{E} [(A_i - \mu)^2 (1 - W_{i,T-1})^2] \leq \mathbb{E} [(A_i - \mu)^2] \mathbb{E} [(1 - W_{i,T-1})^2]. \quad (16)$$

Inequality (16) in Assumption 4.1 is equivalent to

$$\text{Cov} \{ (A_i - \mu)^2, (1 - W_{i,T-1})^2 \} \leq 0.$$

Thus, Assumption 4.1 states that, as two random variables, $(A_i - \mu)^2$ and $(1 - W_{i,T-1})^2$ are weakly negatively correlated. Intuitively, the assumption requires that larger values of $(A_i - \mu)^2$ are associated with smaller weight attributed to the pooled forecast (or that the two are uncorrelated). This is a mild and reasonable assumption, in that it only rules out “pathological” weights that would shrink outliers more than individuals near the mean of the distribution, thus exacerbating the tyranny of the majority phenomenon that we are seeking to overcome. If the individual weight is a fixed constant, i.e., $W_{i,T-1} = c_i$ for some constant $0 \leq c_i \leq 1$, then Assumption 4.1 is satisfied with an equality in (16). Conversely, if the individual weight is a genuine function of the random effect - instead of a fixed constant - the inequality in (16) can be strict. We will see below that this strict inequality translates into improvements in the performance of IW.

⁹See Section A.2 in Manski [2021] and references therein for a detailed discussion.

¹⁰If we had adopted the criterion of minimax instead of minimax regret, what would have mattered between TS and Pool is which $\max_i \sigma_i^2$ or $\max_i \lambda_i^2$ is larger, resulting in an obvious but trivial solution that TS is preferred to Pool if $\max_i \sigma_i^2 < \max_i \lambda_i^2$ and vice versa. In this setting, it is not necessarily the case that IW provides a better performance in terms of minimizing the maximum MSFE.

4.2 Minimax Regret Optimality of IW

In this section, we show the conditions under which IW is optimal in terms of minimax regret. We restrict our attention to the parameter space represented in Figure 1 below, where the signal-to-noise ratio λ_i^2/σ_i^2 ranges from $1 - \nu$ to $1 + \nu$ for some $0 \leq \nu < 1$.

$$\Theta = \Theta(\nu) := \{(\sigma_i^2, \lambda_i^2) \in \mathbb{R}_+^2 : 1 - \nu \leq \lambda_i^2/\sigma_i^2 \leq 1 + \nu\}. \quad (17)$$




Figure 1: Parameter space for $\nu \approx 1$

Considering a neighbourhood of 1 is a natural choice, as we saw that this point represents the case where TS and Pool are equally accurate. The radius of the neighbourhood is constrained by the fact that the signal-to-noise ratio cannot be negative, so in practice the parameter space only excludes cases where Pool performs very poorly relative to TS, due to a large variance of the random effects combined with low variance of the shocks. The adoption of a common ν for the upper and lower bounds is only for convenience in deriving analytical results, and Figures 2 and 3 below make it clear that we are being conservative, as increasing the upper bound on λ_i^2/σ_i^2 well beyond the value of 2 would not change the conclusions of the minimax regret analysis.

The next theorem shows that IW (uniquely) minimizes maximum regret among TS, Pool, and IW under Assumptions 3.1 and 4.1.

Theorem 4.1. *Let Assumptions 3.1 and 4.1 hold. Then,*

$$\max_{\theta_i \in \Theta} R(\text{IW}, \theta_i) \leq \min \left\{ \max_{\theta_i \in \Theta} R(\text{TS}, \theta_i), \max_{\theta_i \in \Theta} R(\text{Pool}, \theta_i) \right\},$$

where Θ is defined in (17). Furthermore, the inequality above is strict if either $0 < W_{i,T-1} < 1$ or the inequality in (16) is strict.

The theorem shows when the improvement of IW over TS and Pool in terms of regret is strict. For example, any constant weight between 0 and 1 will provide an improvement. Furthermore, keeping all else equal, a weight that is a genuine function of the random effect that satisfies Assumption 4.1 with a strict inequality will deliver an additional improvement in performance. Existing Bayesian shrinkage approaches such as the James-Stein estimator, for example, deliver weights that are strictly between 0 and 1 but that do not depend on A_i . This means that James-Stein outperforms (in terms of minimax regret) using TS or Pool for all individuals, but it can be, in turn, outperformed by any weight satisfying Assumption 4.1 that is a genuine function of the random effect. The theorem thus illustrates the potential benefits of individual weights that are based on information in the time series dimension, and thus capture the random effect, relative to existing shrinkage approaches that leverage instead the cross-sectional information.

We illustrate the findings of Theorem 4.1 in Figures 2 and 3. We consider one individual (so we temporarily drop the subscript i) observed over 4 time periods, with shocks U_1, \dots, U_4 drawn independently from a $N(0, 1)$ and random effect A drawn from a $N(0, \lambda^2)$. Repeating the simulation a large number of times allows us to approximate the individual MSFE and regret when forecasting the outcome Y_4 at time $T = 3$ using either TS, Pool or IW. The figures plot these MSFEs and regrets as a function of the signal-to-noise ratio. For IW, we consider the feasible minimax regret optimal rule (IW-MR) that we derive in equations (24) and (25) below.¹¹ Figure 2 shows that no forecast uniformly dominates in terms of MSFE over the parameter space; however, IW is the most accurate over the majority of the parameter space, except for very small values of the signal to noise ratio, when Pool dominates. Figure 3 shows that IW is additionally minimax-regret optimal. To see why, note that the regret for TS (dashed line) achieves a maximum value of around 1 when the signal to noise ratio is close to zero. The regret for Pool (dotted line) obtains a maximum value of around 1.4 when the signal to noise ratio is large. The regret for IW (solid line) achieves a maximum value of around 0.27, when the signal to noise ratio is close to zero. Thus, over the parameter space, the minimax regret optimal rule is IW, since it has the smallest maximum regret among the three rules.

¹¹ Specifically, we have $\hat{Y}_3^{TS} = Y_3$, $\hat{Y}_3^{Pool} = 0$ and $\hat{Y}_3^{IW-MR} = Y_3 W_2$, with $W_2 = 1 - 1/\sqrt{\frac{\max\{Y_1^2, Y_2^2\}}{0.5(Y_1 - Y_2)^2} + 1}$.




Figure 2: MSFE of TS, Pool, and feasible IW




Figure 3: Regret of TS, Pool, and feasible IW

4.3 MSFE Optimality of IW

Figure 2 suggests that IW outperforms TS and Pool in terms of MSFE when the signal to noise ratio is in a neighborhood of 1. One may wonder if this result holds regardless of the data-generating process. The following theorem shows that, indeed, it is possible to show that IW outperforms TS and Pool in terms of MSFE regardless of the data-generating process, if we restrict attention to the case where the signal to noise ratio equals 1. This means that IW is not only robust - i.e. optimal in terms of minimax regret - but it is also optimal in terms of MSFE when TS and Pool are equally accurate and thus would be indistinguishable.

Theorem 4.2. *Let Assumptions 3.1 and 4.1 hold. Suppose that $\lambda_i^2 = \sigma_i^2 = 1$. Then,*

$$\text{MSFE}(\text{IW}, \theta_i) \leq \text{MSFE}(\text{TS}, \theta_i) = \text{MSFE}(\text{Pool}, \theta_i) = 2.$$

Furthermore, the inequality above is strict if either $0 < W_{i,T-1} < 1$ or the inequality in (16) is strict.

The theorem shows that IW is weakly more accurate than TS and Pool when the two forecasts have equal accuracy. Furthermore, as in the case of the minimax regret optimality results in Theorem 4.1, a strict accuracy improvement can be obtained when the weights are strictly between 0 and 1 or are genuine functions of the random effects. This shows that considering individual weights that leverage time series information to capture the random effect can deliver a strict improvement in accuracy.

4.4 Accuracy Gains and Tail Heaviness

In this section we perform a simulation exercise to illustrate the result of Theorem 4.2 and further show how the accuracy gains of IW are linked to the heaviness in the tails of the distribution of A_i . As for Figures 2 and 3, we consider one individual observed over 4 time periods. We now however focus on the case $\sigma^2 = \lambda^2 = 1$ (making TS and Pool equally accurate) with shocks U_1, \dots, U_4 drawn independently from a $N(0, 1)$ and random effect A drawn from a Pareto distribution with different degrees of tail heaviness (in a way that ensures that $\sigma^2 = \lambda^2$).¹² Repeating the simulation a large number of times allows us to approximate the individual MSFE when forecasting the outcome Y_4 at time $T = 3$ using either TS, Pool or IW. For IW we consider the feasible rule that we derive in equations (24) and (25) below (IW-MR),

¹²We consider the double Pareto distribution with pdf $f(x; \theta; \beta) = \theta/(2\beta) \begin{cases} (x/\beta)^{\theta-1}, & \text{if } 0 < x < \beta \\ (\beta/x)^{1-\theta}, & \text{if } x \geq \beta \end{cases}$ with the

following parameter combinations for the shape (θ) and scale (β) parameters: (2.3, .5), (3, 1), (5, 2.45), (50, 34.5). Note that population moments of order θ or greater do not exist; thus, in order to quantify the tail heaviness of the distribution of A_i , we report a robust quantile based measure of kurtosis, the Crow-Siddiqui measure of kurtosis ($CS_k = (Q_{0.975} - Q_{0.025})/(Q_{0.75} - Q_{0.25})$), for each of the above four combinations: 7.58, 6.59, 5.51, 4.42.

which in the simplified setting considered in this section is given in footnote 11. Figure 4 plots these MSFEs of TS, Pool and IW as a function of the heaviness in the tails of the distribution of A , as captured by the Crow-Siddiqui measure of kurtosis (on the x-axis). The figure shows that IW improves on the performance of TS and Pool when TS and Pool are equally accurate, confirming the findings of Theorem 4.2. Furthermore, the accuracy gains of IW are larger the heavier the tails.




Figure 4: MSFE of feasible IFW as a function of tail heaviness

4.5 Validity of Assumption 4.1 and Relationship with Tail Heaviness

It is easy to verify that all the feasible weights reported in Sections 3.2.2, 3.2.1 and 3.2.3 satisfy Assumption 4.1 with a strict inequality. For all weights, the term $(A_i - \mu)^2$ appears in the denominator of $1 - W_{i,T}$, whereas the remaining terms only depend on the $U'_{i,t}$ s. This implies that large values of $(A_i - \mu)^2$ are associated with small values of the weight on Pool. To illustrate how Assumption 4.1 is linked to tail heaviness, we consider the same simulation design as that obtained to produce Figure 4 and compute the covariance in Assumption 4.1 (focusing for simplicity on IW-MR only). The four distributions have increasing tail heaviness, while everything else that would otherwise affect the weights is kept fixed. We find that this covariance becomes more negative as the tail heaviness increases (it respectively equals -0.134, -0.168, -0.216, -0.254 for the four levels of tail heaviness). This implies that heavy tails make the inequality in Assumption 4.1 ‘bite’ more, which, as shown by Theorems 4.2 and 4.1, translates into larger gains of IW relative to TS and Pool.

4.6 Implications for group accuracy

The findings in the previous sections show the benefits of IW in terms of individual performance, but they also have implications for group accuracy. Figure 2, in particular, provides some intuition for how the distribution of the signal to noise ratio λ_i^2/σ_i^2 across i (which we do not restrict in any way) affects group accuracy as measured by the average MSFE. If there are enough individuals for which the signal to noise ratio is in the range where IW dominates all other forecasts in terms of individual MSFE, IW will dominate also in terms of average MSFE across individuals. In addition, our simulations below will show how the tail properties of the distribution of random effects can be linked to improved performance of IW (relative to shrinkage estimators). Another simulation will show that IW can be beneficial not only for outliers, but also for individuals that are near the mean of such distribution. This implies that improvements in terms of group accuracy of our method relative to existing methods can be linked to how many individuals fall in the tails and/or how many are clustered near the mean of the random effects distribution.

5 Feasible Weights for IW

The results in the previous section show that individual weights are optimal under assumption 4.1, but do not directly provide a way to derive feasible weights. In this section, we show that we can derive feasible weights that satisfy this assumption. Here we describe three types of weights. For the first two, as in the previous section, we focus on the simplified setting in section 4.1 for illustrative ease, but, again, the weights that we propose to use in practice and that we consider in the simulations and empirical application are the general weights reported in section 3.2. The third type of weights instead are not based on the model assumptions and, thus, they are directly derived in the general case.

5.1 Estimated Oracle Weights (IW-O)

The first set of feasible weights are based on the oracle weights that minimize the individual MSFE,

$$\text{MSFE}(\hat{Y}_{i,T}^{IW}) = \mathbb{E} \left[\left(Y_{i,T+1} - \hat{Y}_{i,T}^{IW} \right)^2 \right],$$

which are functions of the individual variance parameters:

$$W_i^o = \frac{\lambda_i^2}{\lambda_i^2 + \sigma_i^2}. \quad (18)$$

The oracle weights can be obtained by following Bates and Granger [1969] or Timmermann [2006].¹³ W_i^o could also be obtained by adapting to our context a similar reasoning as in Efron and Morris [1973], who consider a theoretical generalization of James-Stein.¹⁴ The oracle weights depend on the unknown parameters λ_i^2 and σ_i^2 . These parameters are identified in panel data, so one could in principle estimate them using the time series dimension.¹⁵ To our knowledge, feasible individual-specific weights have not been considered in practice in the literature (Efron and Morris [1973] mention the generalized weights as a theoretical possibility in a different context, without discussing feasibility). Estimated oracle weights at time $T - 1$ can be obtained as

$$W_{i,T-1}^{IW-O} = \frac{\sum_{t=1}^{T-1} (Y_{i,t} - \mu)^2 / (T - 1) - \sum_{t=1}^{T-2} (Y_{i,t} - Y_{i,t+1})^2 / 2(T - 2)}{\sum_{t=1}^{T-1} (Y_{i,t} - \mu)^2 / (T - 1)}. \quad (19)$$

These weights use the fact that: $\sum_{t=1}^{T-1} (Y_{i,t} - \mu)^2 / (T - 1)$ is an unbiased estimator of $\lambda_i^2 + \sigma_i^2$ and that $\hat{\sigma}_i^2 = \sum_{t=1}^{T-2} (Y_{i,t} - Y_{i,t+1})^2 / 2(T - 2)$ is an unbiased estimator of σ_i^2 .¹⁶

As one can expect, the short time dimension makes these parameters imprecisely estimated, which is likely to result in poor performance of feasible oracle weights (and the subtraction at the numerator can deliver negative weights that perform very poorly in simulations, see the discussion in Appendix B.1). Our simulations will confirm this intuition and show that the alternative feasible weights we discuss in the next section are preferable.

The above considerations on the likelihood of poor performance of feasible oracle weights lead us to additionally focus on obtaining feasible weights that are robust. Specifically, we seek to obtain weights that are minimax regret optimal over the unknown parameter space.

¹³The linear combination (14) and the weight (18) follow from equation (9) in Chapter 4 of Timmermann [2006], using the fact that the joint distribution of $Y_{i,T+1}$ and $Y_{i,T}$ is

$$\begin{pmatrix} Y_{i,T+1} \\ Y_{i,T} \end{pmatrix} \sim \left(\begin{pmatrix} \mu \\ \mu \end{pmatrix}, \begin{pmatrix} \lambda_i^2 + \sigma_i^2 & \lambda_i^2 \\ \lambda_i^2 & \lambda_i^2 + \sigma_i^2 \end{pmatrix} \right)$$

which gives the optimal weight on $\hat{Y}_{i,T}^{TS}$ as the product between the inverse of the variance of the forecast and the covariance between the outcome and the forecast, i.e.: $W_i^o = \frac{\lambda_i^2}{\lambda_i^2 + \sigma_i^2}$.

¹⁴The “best linear rule” in equation (9.4), page 129 of Efron and Morris [1973] coincides with our equation (14) with weight (18) when applied to $Y_{i,T}|A_i \sim (A_i, \sigma_i^2)$ and $A_i \sim (\mu, \lambda_i^2)$.

¹⁵Note that existing shrinkage estimators obtain feasible weights under the assumption of homogeneous λ_i^2 and σ_i^2 (James and Stein [1961]) or homogeneous λ_i^2 (Kwon [2021]), which can then be estimated using pooled estimators that leverage the cross-sectional dimension.

¹⁶To see that $\hat{\sigma}_i^2$ is an unbiased estimator of σ_i^2 note that:

$$\mathbb{E} \left[\sum_{t=1}^{T-2} (Y_{i,t} - Y_{i,t+1})^2 \right] = \mathbb{E} \left[\sum_{t=1}^{T-2} U_{i,t}^2 + \sum_{t=1}^{T-2} U_{i,t+1}^2 \right] = 2(T - 2)\sigma_i^2.$$

5.2 Minimax-Regret Optimal Weights (IW-MR)

In order to obtain feasible minimax-regret optimal weights we shift from unconditional MSFE to MSFE that is conditional on the information set at time $T - 1$. The following lemma is the analog of Lemma 4.1 for the conditional MSFE.

Lemma 5.1. *Consider the forecasts in (12), (13), and (14). Let Assumption 3.1 hold. Then, the mean squared forecast errors conditional on the information set at time $T-1$, $\{Y_{i,1}, \dots, Y_{i,T-1}\}$ are*

$$\begin{aligned} \text{MSFE}(\text{TS}, \theta_i | \{Y_{i,1}, \dots, Y_{i,T-1}\}) &= 2\sigma_i^2, \\ \text{MSFE}(\text{Pool}, \theta_i | Y_{i,1}, \dots, Y_{i,T-1}) &= \kappa_{i,T-1}^2 + \sigma_i^2, \\ \text{MSFE}(\text{IW}, \theta_i | Y_{i,1}, \dots, Y_{i,T-1}) &= \sigma_i^2(1 + W_{i,T-1}^2) + \kappa_{i,T-1}^2(1 - W_{i,T-1})^2, \end{aligned}$$

where

$$\kappa_{i,T-1}^2 := \mathbb{E}[(A_i - \mu)^2 | Y_{i,1}, \dots, Y_{i,T-1}]. \quad (20)$$

It is easy to verify that the weights that minimize $\text{MSFE}(\text{IW}, \theta_i | Y_{i,1}, \dots, Y_{i,T-1})$ are given by:

$$W_{i,T-1}^* = \kappa_{i,T-1}^2 / (\kappa_{i,T-1}^2 + \sigma_i^2). \quad (21)$$

In this section we consider a different type of regret, defined as the difference between the conditional MSFE for a generic weight $W_{i,T-1}$ and the conditional MSFE that corresponds to the conditionally optimal weights $W_{i,T-1}^*$ in (21):

$$\begin{aligned} R^*(W_{i,T-1}, \theta_i | Y_{i,1}, \dots, Y_{i,T-1}) &:= \text{MSFE}(W_{i,T-1}, \theta_i | Y_{i,1}, \dots, Y_{i,T-1}) - \text{MSFE}(W_{i,T-1}^*, \theta_i | Y_{i,1}, \dots, Y_{i,T-1}) \\ &= \sigma_i^2 W_{i,T-1}^2 + \kappa_{i,T-1}^2 (1 - W_{i,T-1})^2 - \frac{\sigma_i^2 \kappa_{i,T-1}^2}{\kappa_{i,T-1}^2 + \sigma_i^2} \\ &= \sigma_i^2 \left[W_{i,T-1}^2 + \zeta_{i,T-1}^2 (1 - W_{i,T-1})^2 - \frac{\zeta_{i,T-1}^2}{\zeta_{i,T-1}^2 + 1} \right], \end{aligned} \quad (22)$$

where

$$\zeta_{i,T-1}^2 := \frac{\kappa_{i,T-1}^2}{\sigma_i^2} = \frac{\mathbb{E}[(A_i - \mu)^2 | Y_{i,1}, \dots, Y_{i,T-1}]}{\sigma_i^2}. \quad (23)$$

The form of regret in (22) is similar to that of regret in statistical decision theory (e.g., see equation (6) in Manski [2021]).

The following theorem obtains the optimal minimax regret weights under the assumption

that we can put bounds on the random variable $\zeta_{i,T-1}^2$.

Theorem 5.1. *Let Assumption 3.1 hold. Suppose that it is known that $\zeta_{i,T-1}^2$ in (23) is such that $\zeta_{i,T-1}^2 \in [0, \tilde{\zeta}_{i,T-1}^2]$, where $\tilde{\zeta}_{i,T-1}^2$ is large enough that maximum regret can occur at $\zeta_{i,T-1}^2 = \tilde{\zeta}_{i,T-1}^2$. Consider maximum regret*

$$\max_{\theta_i \in \Theta} R^*(W_{i,T-1}, \theta_i | Y_{i,1}, \dots, Y_{i,T-1}) = \sigma_i^2 \max \left[W_{i,T-1}^2, \left\{ W_{i,T-1}^2 + \tilde{\zeta}_{i,T-1}^2 (1 - W_{i,T-1})^2 - \frac{\tilde{\zeta}_{i,T-1}^2}{\tilde{\zeta}_{i,T-1}^2 + 1} \right\} \right],$$

with $R^*(W_{i,T-1}, \theta_i | Y_{i,1}, \dots, Y_{i,T-1})$ defined as in (22). Then, the weight that minimizes maximum regret is

$$W_{i,T-1}^{IW-MR} = 1 - \frac{1}{\sqrt{\tilde{\zeta}_{i,T-1}^2 + 1}}. \quad (24)$$

In practice, the value of the bound $\tilde{\zeta}_{i,T-1}^2$ is uncertain, but the following heuristic rule can be used to obtain feasible weights: assuming $T \geq 3$,

$$\widehat{\tilde{\zeta}_{i,T-1}^2} := \frac{\max\{(Y_{i,1} - \mu)^2, \dots, (Y_{i,T-1} - \mu)^2\}}{\sum_{t=1}^{T-2} (Y_{i,t} - Y_{i,t+1})^2 / 2(T-2)}, \quad (25)$$

where, again, μ is either known or approximated by the pooled mean. Intuitively, the denominator $\sum_{t=1}^{T-2} (Y_{i,t} - Y_{i,t+1})^2 / 2(T-2)$ is an unbiased estimator of the denominator of $\zeta_{i,T-1}^2$, σ_i^2 . The numerator is a proxy for the upper bound on $\kappa_{i,T-1}^2$ in (20), which is the numerator of $\zeta_{i,T-1}^2$. One could also consider alternative proxies.¹⁷

5.3 Inverse MSFE Weights (IW-MSFE)

The weights we derive in this section do not rely on the model assumptions, and are thus applicable in the general setting considered in section 2. The idea is to compute individual weights by comparing the (in-sample or out-of-sample) MSFE at time T of the competing forecasts. These weights are analogous to those considered in the time-series forecast combination literature (e.g., Stock and Watson [1998]), with the difference that the MSFE is computed here for each individual over a very small time series sample (possibly containing only one observation). As in Stock and Watson [1998], these inverse MSFE weights ignore any correlation between the forecasts in the combination.¹⁸

¹⁷For example, one could multiply (25) by $\frac{T-1}{T}$ based on the fact that, if $\zeta_{i,T-1}^2$ in (23) is uniformly distributed in $[0, \tilde{\zeta}_{i,T-1}^2]$, $\hat{M} = \max_{1 \leq t \leq T-1} (\zeta_{i,t}^2)$ has bias $\mathbb{E}[\hat{M}] = \int_0^{\tilde{\zeta}_{i,T-1}^2} \frac{T-1}{(\tilde{\zeta}_{i,T-1}^2)^{T-1}} t^{T-1} dt = \frac{T-1}{T} \tilde{\zeta}_{i,T-1}^2$. In practice, we do not recommend doing so, given our focus on a distribution-free environment.

¹⁸In the time series literature, these weights are known to perform well even when the time dimension is large because of the challenges in estimating correlations precisely. See, e.g., the discussion in Stock and Watson [1998].

The in-sample inverse MSFE weights are given by:

$$W_{i,T}^{IW-MSFE-IS} := \frac{1/\left[\sum_{t=1}^T (Y_{i,t} - \hat{Y}_{i,T}^{TS})^2\right]}{1/\left[\sum_{t=1}^T (Y_{i,t} - \hat{Y}_{i,T}^{TS})^2\right] + 1/\left[\sum_{t=1}^T (Y_{i,t} - \hat{Y}_{i,T}^{Pool})^2\right]}. \quad (26)$$

The out-of-sample inverse MSFE weights are given by:

$$W_{i,T}^{IW-MSFE-OOS} := \frac{(Y_{i,T} - \hat{Y}_{i,T-1}^{TS})^{-2}}{(Y_{i,T} - \hat{Y}_{i,T-1}^{TS})^{-2} + (Y_{i,T} - \hat{Y}_{i,T-1}^{Pool})^{-2}}. \quad (27)$$

Note that here we base $W_{i,T}^{OOS}$ only on the out-of-sample forecast errors at time T corresponding to TS and Pool forecasts computed on the sample up to time $T-1$. Depending on the magnitude of T , one could also compute the out-of-sample MSFEs using more than just one out-of-sample period. For example, one could select $P < T$ and consider

$$W_{i,T,P}^{IW-MSFE-OOS} := \frac{1/\left[\sum_{t=T-P+1}^T (Y_{i,t} - \hat{Y}_{i,t-1}^{TS})^2\right]}{1/\left[\sum_{t=T-P+1}^T (Y_{i,t} - \hat{Y}_{i,t-1}^{TS})^2\right] + 1/\left[\sum_{t=T-P+1}^T (Y_{i,t} - \hat{Y}_{i,t-1}^{Pool})^2\right]}. \quad (28)$$

Finally, we note that one could consider “rolling-window” forecasts, both as the original TS and Pool forecasts and in the computation of the weights. In this case, both TS and Pool forecasts at time t would be based only on the $R < t$ most recent observations, rather than all available observations up to time t .

6 Monte Carlo Simulations

In this section we first study the finite sample performance of alternative feasible IW weights. Then, we compare IW to existing Bayesian shrinkage.

6.1 Alternative Feasible IW

We consider one individual (so we here drop the subscript i) observed over 3 time periods, with shocks U_1, \dots, U_3 drawn independently from a $N(0, 1)$ and random effect A drawn from a $N(0, \lambda_d^2)$, with λ_d^2 taking $d = 50$ equally-spaced values on the grid $[0.001, 2]$. Repeating the simulation 10000 times allows us to approximate the individual MSFE when forecasting the outcome Y_3 at time $T = 2$ using the various different feasible weights for IW considered in Section 3.2 (setting $\hat{Y}_{i,T}^{Pool} = 0$). Figures 5 and 6 report the various MSFEs, divided by the MSFE of IW-MR, and the minimax regret of different feasible IW rules as a function of the signal-to-noise ratio (which here equals λ^2). In Figure 5, a line above 1 means that the rule is uniformly dominated by IW-MR over the parameter space.




Figure 5: MSFE of alternative feasible IW rules.




Figure 6: Regret of alternative feasible IW rules.

Figures 5 and 6 illustrate the dominance of IW-MR (black solid line) over the other feasible rules, in terms of robustness over the parameter space. IW-MR is minimax regret optimal over the considered parameter space. IW-MSFE-IS (blue dotted line) is uniformly dominated by IW-MR, although not by a large amount. The performance of IW-MSFE-OOS (red dashed line) and IW-O (green dashed-dotted line) is similar and depends on the signal to noise ratio, outperforming IW-MR when the signal to noise ratio is very low, but performing poorly over the rest of the parameter space.

6.2 IW vs. Bayesian Shrinkage

In this section we compare the performance of IW-MR to that of the James Stein forecast. The draws in the following simulations can be interpreted in two different ways. First, they can be seen as different possible draws of the random effect for one individual. Second, they can be seen as draws for different individuals that have the same distribution of the random effect. Averages across simulations accordingly have a different interpretation (e.g., they approximate the individual MSFE in the first interpretation and the group MSFE in the second). The James-Stein forecast in the first interpretation is just an individual weighting rule with constant weights that do not depend on the individual random effect (see the discussion after Theorem 4.1). In the second interpretation, the James-Stein is the forecast that exploits information from the cross-sectional dimension, in contrast to IW, which leverages the time series dimension. Note that the assumptions of parameter homogeneity made by James-Stein are satisfied in all the designs considered below.

6.2.1 Tyranny of the Majority

We start by visually illustrating how IW overcomes the “tyranny of the majority” phenomenon that affects existing Bayesian shrinkage methods. Henceforth, we focus on the IW-MR rule, which we saw in the previous section generally outperforms the other feasible rules. We consider 10000 simulations of outcomes generated as $Y_t = A + U_t$, with $t = 1, \dots, 3$, $U_t \sim N(0, 1)$, independent across t . For the random effects we consider the following designs:

- Design 1 (Normal): $A \sim N(0, \lambda_d^2)$, with λ_d taking $d = 2$ values: $[1, 3]$.
- Design 2 (Laplace): $A \sim \text{Laplace}$ with parameters $(0, 1)$, which implies mean 0 and variance $\lambda^2 = 2$.
- Design 3 (Double Pareto): $A \sim \text{DoublePareto}(\theta, \beta)$, where $\theta = 3$ and $\beta = 1$ (which implies mean 0 and variance λ^2 around 1.1).

These designs correspond to an increasing heaviness in the tails of the distribution of the random effect A .

We compare IW-MR as described in Section 3.2 (with $\hat{Y}_T^{Pool} = 0$) to the James-Stein forecast (JS):

$$\hat{Y}_T^{JS} := \frac{\lambda^2}{\lambda^2 + \sigma^2/T} \hat{Y}_T^{TS}, \quad (29)$$

where $\sigma^2 = 1$ in all designs and λ^2 is indicated in each design.

Figures 7 - 10 report the difference ΔSFE between the squared forecast errors of forecasts made at time $T = 2$ for IW-MR and those for JS, for the different designs. The horizontal axis reports the value of A . The figures illustrate the tyranny of the majority phenomenon: JS tends to make larger errors than IW-MR (the dots fall below zero) when the random effects fall in the tails and also near the center of the distribution. This pattern is not yet visible in Figure 7 for the normal design with low variance, where the cloud appears symmetric relative to the horizontal axis, but it is clear in the remaining figures. For example, in Figure 8 (the normal design with larger variance) the cloud is heart-shaped, showing the superior performance of IW-MR near the mean of the distribution. Figure 9 (the Laplace design) also shows the heart shape but also the superior performance of IW in the tails. The improvement in the tails is starkly evident in Figure 10 (the Double Pareto design), where the cloud has an inverted U-shape. These results illustrate that what matters for the tyranny of the majority is not only the tail heaviness of the distribution of random effects, but its relationship to the variance: the design in Figure 8 shows that the phenomenon is present even when the distribution has thin tails, but has a large variance. This is intuitive, as both a high variance and heavy tails make it worthwhile to link the shrinkage to where the individual falls in the distribution of random effects (IW) instead of shrinking every individual by the same amount (JS).




Figure 7: Tyranny of the majority.




Figure 8: Tyranny of the majority.




Figure 9: Tyranny of the majority.




Figure 10: Tyranny of the majority.

6.2.2 Group Accuracy

The second interpretation of the simulation discussed above, which views the different draws as different individuals, allows us to also analyze the group accuracy in the different designs. Averaging the ΔSFE 's reported in each figure gives a measure of the relative group accuracy of IW-MR and JS. This quantity equals 0.019, 0.025, -0.005 and -0.027, respectively in Figures 7 - 10. The relative group accuracy of JS and IW-MR for these designs thus depends on the tail properties of the distribution of random effects, with JS dominating in the normal cases and IW-MR dominating in the heavy-tailed cases.

7 Empirical Applications

We consider two applications of IW to estimation of random effects and to microforecasting.

7.1 Estimating and Forecasting Systemic Firm Discrimination

In this section we use IW to replicate and extend the analysis in Kline et al. [2022], assessing the extent to which large U.S. employers systemically discriminate job applicants based on gender.

7.1.1 Data

We use the panel dataset in Kline et al. [2022] on an experiment that consisted of sending fictitious applications to jobs posted by 108 of the largest U.S. employers. For each firm, 125 entry-level vacancies were sampled and, for each vacancy, 8 job applications with random characteristics were sent to the employer. Sampling was organized in 5 waves (between October

2019 and April 2021). Focusing on firms sampled in all waves yields a balanced panel of $N = 72$ firms over $T = 5$ waves.¹⁹

Applications were sent in pairs, one randomly assigned a distinctively female name and the other a distinctively male name. For details on the other observables see Kline et al. [2022].

The primary outcome in Kline et al. [2022] is whether the employer attempted to contact the applicant within 30 days of applying. The gender contact gap is defined as the firm-level difference between the contact rate (the ratio of number of contacts and number of received applications) for male and that for female applications.

Kline et al. [2022] estimate the cross-firm distribution of discrimination using the empirical Bayes deconvolution estimator of Efron [2016], relying on a normality assumptions for the idiosyncratic shocks. The estimator considers firm-specific studentized contact gaps, defined as $y_{i,t} = Y_{i,t}/s_i$, where $Y_{i,t}$ is the contact gap and s_i is the standard deviation of contact gaps across different job applications for firm i . These are modelled as

$$y_{i,t} = a_i + u_{i,t}, \quad u_{i,t} \sim N(0, 1) \quad a_i \sim G_a, \quad \text{for } i = 1, \dots, 72.$$

The distribution G_a of the random effect is assumed to belong to an exponential family, flexibly parameterized by a fifth-order spline. By pooling observations from all five waves, Efron [2016]’s approach yields penalized Maximum Likelihood estimates of the spline parameters and thus an implied distribution \hat{G}_a of studentized contact gaps with corresponding density $\hat{g}_a = d\hat{G}_a$. One can then recover the distribution \hat{G}_Y of the random effects for the unstudentized contact gaps Y_{it} under the assumption of independence between the random effects and s_i . In particular, the density $\hat{g}_Y = d\hat{G}_Y$ at each point x is obtained as $\hat{g}_Y(x) = \frac{1}{N} \sum_{i=1}^N \frac{1}{s_i} \hat{g}_a\left(\frac{x}{s_i}\right)$.²⁰

7.1.2 Estimating the Distribution of Firm Discrimination via IW

In this section, we estimate the cross-firm distribution of discrimination implied by IW, and compare it to the estimate obtained by applying Efron [2016], as in Kline et al. [2022]. Because IW does not rely on distributional assumptions, the findings can help assess the validity of the normality assumption on the idiosyncratic shocks imposed by Kline et al. [2022].

For each $t = 1, \dots, 5$ and firm i we model the unstudentized gender contact gaps $Y_{i,t}$ as:

$$Y_{i,t} = A_i + U_{i,t}. \tag{30}$$

At $T = 5$, we produce individual one-step-ahead forecasts using IW-MR as discussed in section

¹⁹Accounting for vacancy closures and the exclusion of some firms from some waves reduced the number of applications to 65,400.

²⁰To perform the deconvolution, the choice of two tuning parameters is required: the order of the spline and the penalization parameter of the first-step maximum likelihood procedure. The latter is optimally calibrated to obtain a variance matching the bias-corrected estimate in Table IV of Kline et al. [2022].

3.2. The distribution of these forecasts across firms can be viewed as an estimate of firm discrimination, based on the information from all five waves. Figure 11 plots a non-parametric estimate of the density.²¹ The figure also plots the density obtained by applying Efron [2016]’s approach, also based on the information from all five waves.




Figure 11: Density of gender contact gaps estimated using IW and Efron [2016]

We find that IW and Efron [2016] deliver densities that are visually not too different, but the IW-implied density is slightly more right-skewed and has a heavier right tail. As a measure of the overall inequality, we further compute the Gini coefficient as in Section IX of Kline et al. [2022] and obtain similar coefficients of 0.489 using IW and 0.496 using Efron [2016]. For the purpose of quantifying inequality, therefore, the two approaches deliver similar conclusions, however IW suggests that a larger proportion of firms discriminate in favour of male applicants (61% for IW vs. 54% for Efron [2016]).

7.1.3 Forecasting Performance

We then consider the forecasting performance of IW relative to that of Efron [2016]’s (henceforth Efron) and James-Stein’s (henceforth JS) approaches. For each wave $T = 3, 4$, we produce individual one-step-ahead forecasts of contact gaps by the following methods: TS, which uses the time-series mean at time T ; Pool, which uses the pooled mean at time T ; the feasible rule IW-MR from Section 3.2 and JS as in equation (29). For Efron [2016], we obtain individual forecasts by the Empirical Bayes posterior mean estimates.²² We then compare the individual

²¹This is obtained by applying the default option in the “density” function from the Stats package in R, which uses a kernel density estimator.

²²We use the code provided by Kline et al. [2022] that produces Figure A13 in their paper, where they assess the out-of-sample predictive power of the posterior means. We adapt the code to use data from waves $1, \dots, T$

out-of-sample forecasts from each method k , $\{\hat{y}_{i,T}^k\}$ to the actual realizations $\{y_{i,T+1}\}$, for the waves 4, 5.

For each forecasting method k and each firm i , the mean squared forecast error over the out-of-sample period is

$$MSFE(k, i) = \frac{1}{2} \sum_{T=3}^4 (y_{i,T+1} - \hat{y}_{i,T}^k)^2.$$

We report the average MSFE across i in Table 1.

Table 1: Out-of-sample Accuracy - Average MSFE

TS	Pool	JS	Efron	IW-MR
.00297	.00336	.00312	.00314	.00294

The table reveals that the IW-MR rules outperform Efron and JS (as well as TS and Pool) in terms of group forecasting accuracy.

Figure 12 reports the difference ΔSFE between the squared forecast errors of forecasts for IW-MR and those for Efron. The horizontal axis reports the value of the gender contact gap at $T = 4$. The figure reveals that Efron - besides being less accurate than IW-MR on average - tends to make larger errors (the dots fall below zero) for firms that at the time of forecasting fell in the right tail and near the center of the distribution. Overall, these findings could reflect a violation of the normality assumption and/or the effect of the data-dependent choice of regularization parameter used by Kline et al. [2022] when implementing Efron.




Figure 12: Difference in Squared Forecast Errors between IW-MR and Efron.

to produce the forecast at $T = 3, 4$.

7.2 Microforecasting Earnings

This section considers an out-of-sample exercise that applies IW to forecasting earnings residuals using the Panel Study of Income Dynamics (PSID).

7.2.1 Data

We consider earnings data from the PSID for 1968-1993.²³ We follow the literature on income dynamics (e.g., Meghir and Pistaferri [2004]) and select a sample of male workers, heads of household, aged between 24 and 55 (inclusive). We drop individuals identifying as Latino, with a spell of self-employment, with zero or top-coded wages and with missing records on race and education. We also require that the change in log earnings is not greater than +5 or less than -3. We consider earnings residuals obtained from a first stage panel data regression of log labor income of an individual i at time t , $\tilde{Y}_{i,t}$, on education, a quadratic polynomial in age, race and year dummies. We denote by $Y_{i,t}$ the residuals from this regression. The goal is to obtain individual one-year-ahead forecasts of earnings residuals $Y_{i,t}$.²⁴

7.2.2 Forecasting Performance

We compare the out-of-sample group accuracy of IW-MR from Section 3.2, versus using TS or Pool for all individuals.

We report results for the balanced samples of $N = 164$ ($N = 790$) individuals with continuous earnings in all consecutive years for 1968-1993 (1968-1980). We further consider an unbalanced sample built using rolling windows of $T = 3$ time periods of balanced samples of individuals (which delivers sample sizes ranging from 3960 to 7912). Forecasts are based on the model:

$$Y_{i,t} = A_i + U_{i,t}. \quad (31)$$

We use rolling windows of $T = 2$ time periods and compare the out-of-sample forecasts from each method k , $\hat{y}_{i,T}^k$, to the actual realizations $y_{i,T+1}$, for $t = 1972, \dots, 1992$, $i = 1, \dots, N$.

For each forecasting method k and each individual i , the mean squared forecast error over the out-of-sample period is

²³We use data up to 1993 because from 1994 a major revision of the survey disrupted the continuity of PSID files, see Kim, Loup, Lupton, and Stafford [2000]. Moreover, after 1997 the PSID switched from an annual to a biannual data collection.

²⁴Forecasting earnings residuals is of interest since they measure individual income risk. For instance, accurate forecasting of individual earnings residuals might be useful for prospective lenders when deciding on loan applications.

$$MSFE(k, i) = \frac{1}{21} \sum_{T=1972}^{1992} \frac{1}{N} \sum_{i=1}^N (y_{i,T+1} - \hat{y}_{i,T}^k)^2. \quad (32)$$

Table 1 reports averages of $MSFE(k, i)$ across i for each forecasting method k .

Table 2: Average out-of-sample MSFE

Sample Size N	TS	Pool	IW-MR
164	0.075	0.211	0.070
794	0.069	0.220	0.067
Unbal. 4-8000	0.117	0.265	0.108

Table 2 shows that, while TS clearly outperforms Pool in terms of average MSFE, IW further improves accuracy.

To gain some insight into which individuals are given higher weights to pooling by IW, in Figure 13 we divide the N=164 individuals of the balanced sample into ten quantiles according to their lagged earnings (the vertical axis) for each year (the horizontal axis). Within each quantile we compute the forecasts given the higher weights by IW-MR: the size of the dots is proportional to the average weight attributed by IW-MR to Pool across individuals in that year and in that quantile.²⁵ Figure 13 shows that individuals near the median of the earnings residuals distribution are those who benefit from larger weights to pooling.




Figure 13: Average weights attributed to Pool by IW-MR by year and earnings quantiles

One possible interpretation of our findings is that in the PSID there is enough unobserved

²⁵We set the size option of the R package ggplot equal to the mean of weights attributed to Pool by the IW-MR rule for each quantile.

heterogeneity to make the time series forecast outperform pooling (as indicated by Table 2). However, an additional improvement in accuracy can be obtained by using IW, which tends to attribute higher weight to pooling for individuals near the median of the distribution. This finding confirms the usefulness of IW even when there is only a small fractions of individuals that would benefit from pooling.

8 Conclusion

Estimating random effects and forecasting with micropanel data is challenging due to the short time dimension, and existing solutions have drawbacks. We presented an alternative method that overcomes these drawbacks, while requiring minimal assumptions. In practice, the method shrinks the time series mean towards the panel mean, using weights that are individual-specific and are computed only using time series information. We offer three types of feasible weights: minimax-regret optimal weights, estimated oracle weights and inverse MSFE weights, with minimax-regret optimal weights generally displaying the best performance.

The method applies to linear panel data models and value-added models, provided the covariates have homogeneous coefficients. The extension to models with heterogeneous coefficients for covariates (i.e., to multivariate unobserved heterogeneity) and to loss functions different from the squared error considered here is straightforward for two of the feasible weights that we consider. In particular, the inverse-MSFE error weights are directly applicable to models with multivariate unobserved heterogeneity and to different loss functions. The estimated oracle weights are also easily extended to multivariate unobserved heterogeneity by modifying the analysis in Pesaran et al. [2022] to allow for heterogeneous variance of the parameters, which in practice means estimating the weights that they propose using time series- instead of cross-sectional information. Extending the derivation of minimax-regret optimal weights to more general settings is less straightforward, and we thus leave this endeavour for future research.

A Proofs

Proof of Lemma 4.1. The MSFEs for TS and Pool are immediate. For IW, first write

$$Y_{i,T+1} - \hat{Y}_{i,T}^{IW} = (Y_{i,T+1} - Y_{i,T}) W_{i,T-1} + (Y_{i,T+1} - \mu)(1 - W_{i,T-1}).$$

Then, under assumption 3.1, we have that

$$\begin{aligned} & \text{MSFE}(\text{IW}, \theta_i) \\ &= \mathbb{E} \left[\left(Y_{i,T+1} - \hat{Y}_{i,T}^{IW} \right)^2 \right] \\ &= \mathbb{E} \left[(Y_{i,T+1} - Y_{i,T})^2 \right] \mathbb{E} \left[W_{i,T-1}^2 \right] + \mathbb{E} \left[(Y_{i,T+1} - \mu)^2 (1 - W_{i,T-1})^2 \right] \\ &\quad + 2\mathbb{E} \left[(Y_{i,T+1} - Y_{i,T})(Y_{i,T+1} - \mu) W_{i,T-1} (1 - W_{i,T-1}) \right] \\ &= 2\sigma_i^2 \mathbb{E} \left[W_{i,T-1}^2 \right] + \mathbb{E} \left[(A_i - \mu)^2 (1 - W_{i,T-1})^2 \right] + \sigma_i^2 \mathbb{E} \left[(1 - W_{i,T-1})^2 \right] \\ &\quad + 2\mathbb{E} \left[(U_{i,T+1})^2 W_{i,T-1} (1 - W_{i,T-1}) \right] \\ &= 2\sigma_i^2 \mathbb{E} \left[W_{i,T-1}^2 \right] + \mathbb{E} \left[(A_i - \mu)^2 (1 - W_{i,T-1})^2 \right] + \sigma_i^2 \mathbb{E} \left[(1 - W_{i,T-1})^2 \right] \\ &\quad + 2\sigma_i^2 \mathbb{E} \left[W_{i,T-1} (1 - W_{i,T-1}) \right] \\ &= \sigma_i^2 + \sigma_i^2 \mathbb{E} \left[W_{i,T-1}^2 \right] + \mathbb{E} \left[(A_i - \mu)^2 (1 - W_{i,T-1})^2 \right], \end{aligned}$$

which proves the lemma. □

Lemma A.1 is used to prove Theorem 4.1.

Lemma A.1. *Let $\mathcal{M} = \{\text{TS}, \text{Pool}, \text{IW}\}$. Let Assumptions 3.1 and 4.1 hold. Then,*

$$R(\text{IW}, \theta_i) \leq \sigma_i^2 \nu$$

for each $\theta_i \in \Theta$, which is defined in (17). Furthermore, the inequality above is strict if either $0 < W_{i,T-1} < 1$ or the inequality in (16) is strict.

Proof of Lemma A.1. To bound $\text{MSFE}(\text{IW}, \theta_i)$, invoke Assumption 4.1 to write

$$\mathbb{E} \left[(A_i - \mu)^2 (1 - W_{i,T-1})^2 \right] \leq \lambda_i^2 \mathbb{E} \left[(1 - W_{i,T-1})^2 \right].$$

This implies that

$$\text{MSFE}(\text{IW}, \theta_i) \leq \sigma_i^2 + \sigma_i^2 \mathbb{E} \left[W_{i,T-1}^2 \right] + \lambda_i^2 \mathbb{E} \left[(1 - W_{i,T-1})^2 \right]. \quad (33)$$

Note that

$$R(\text{IW}, \theta_i) = \max \left\{ 0, \text{MSFE}(\text{IW}, \theta_i) - \sigma_i^2 - \min\{\sigma_i^2, \lambda_i^2\} \right\}.$$

If $\text{MSFE}(\text{IW}, \theta_i) < \sigma_i^2 + \min\{\sigma_i^2, \lambda_i^2\}$, then $R(\text{IW}, \theta_i) = 0$. In this case, there is nothing left to prove. Hence, it suffices to assume that $\text{MSFE}(\text{IW}, \theta_i) \geq \sigma_i^2 + \min\{\sigma_i^2, \lambda_i^2\}$. It follows from (33) and Assumption 4.1 that

$$\begin{aligned}
& \text{MSFE}(\text{IW}, \theta_i) - \sigma_i^2 - \min\{\sigma_i^2, \lambda_i^2\} \\
& \leq \sigma_i^2 \mathbb{E}[W_{i,T-1}^2] + \lambda_i^2 \mathbb{E}[(1 - W_{i,T-1})^2] - \min\{\sigma_i^2, \lambda_i^2\} \\
& \leq \max\{\sigma_i^2, \lambda_i^2\} (\mathbb{E}[W_{i,T-1}^2] + \mathbb{E}[(1 - W_{i,T-1})^2]) - \min\{\sigma_i^2, \lambda_i^2\} \\
& \leq \max\{\sigma_i^2, \lambda_i^2\} - \min\{\sigma_i^2, \lambda_i^2\} \\
& = (\lambda_i^2 - \sigma_i^2) \mathbb{I}(\lambda_i^2 > \sigma_i^2) + (\sigma_i^2 - \lambda_i^2) \mathbb{I}(\lambda_i^2 < \sigma_i^2) \\
& \leq \sigma_i^2 \nu,
\end{aligned} \tag{34}$$

where $\mathbb{I}\{\cdot\}$ denotes the indicator variable and the third inequality uses the fact that $W_{i,T-1}^2 + (1 - W_{i,T-1})^2 \leq 1$ if $0 \leq W_{i,T-1} \leq 1$. In conclusion, we have shown that $R(\text{IW}, \theta_i) \leq \sigma_i^2 \nu$ for each $\theta_i \in \Theta$. This proves the first conclusion of the lemma. The second conclusion follows from the facts that the inequality in (33) will be strict if the inequality in (16) is strict and that the third inequality in (34) will be strict if $0 < W_{i,T-1} < 1$. \square

Proof of Theorem 4.1. We have

$$\min_{m \in \mathcal{M}} \text{MSFE}(m, \theta_i) \leq \min_{m \in \{\text{TS}, \text{Pool}\}} \text{MSFE}(m, \theta_i) = \sigma_i^2 + \min\{\sigma_i^2, \lambda_i^2\}.$$

Furthermore, the regrets for TS and Pool are

$$\begin{aligned}
R(\text{TS}, \theta_i) & \geq \sigma_i^2 - \min\{\sigma_i^2, \lambda_i^2\}, \\
R(\text{Pool}, \theta_i) & \geq \lambda_i^2 - \min\{\sigma_i^2, \lambda_i^2\}.
\end{aligned}$$

Note that

$$\begin{aligned}
\max_{\theta_i \in \Theta} R(\text{TS}, \theta_i) & \geq \max_{\theta_i \in \Theta} [(\sigma_i^2 - \lambda_i^2) \mathbb{I}\{\sigma_i^2 > \lambda_i^2\}] = \sigma_i^2 \nu, \\
\max_{\theta_i \in \Theta} R(\text{Pool}, \theta_i) & \geq \max_{\theta_i \in \Theta} [(\lambda_i^2 - \sigma_i^2) \mathbb{I}\{\sigma_i^2 < \lambda_i^2\}] = \sigma_i^2 \nu,
\end{aligned} \tag{35}$$

where $\mathbb{I}\{\cdot\}$ denotes the indicator variable as before. The claim in Theorem 4.1 then follows directly from Lemma A.1 and the inequalities in (35). \square

Proof of Theorem 4.2. If $\lambda_i^2 = \sigma_i^2 = 1$, it follows from (34) that

$$\text{MSFE}(\text{IW}, \theta_i) - 2 \leq 0,$$

which proves the first conclusion of the theorem. As in Lemma A.1, the second conclusion follows from the facts that the inequality in (33) will be strict if the inequality in (16) is strict

and that the third inequality in (34) will be strict if $0 < W_{i,T-1} < 1$. \square

Proof of Lemma 5.1. First, consider the MSFE for the TS forecast. Write

$$\begin{aligned}\mathbb{E}[(Y_{i,T+1} - Y_{i,T})^2 | Y_{i,1}, \dots, Y_{i,T-1}] &= \mathbb{E}[(U_{i,T+1} - U_{i,T})^2 | Y_{i,1}, \dots, Y_{i,T-1}] \\ &= \mathbb{E}[(U_{i,T+1} - U_{i,T})^2] = 2\sigma_i^2.\end{aligned}$$

Now consider the MSFE for the Pool forecast. Note that

$$\begin{aligned}\mathbb{E}[(Y_{i,T+1} - \mu)^2 | Y_{i,1}, \dots, Y_{i,T-1}] &= \mathbb{E}[(A_i - \mu + U_{i,T+1})^2 | Y_{i,1}, \dots, Y_{i,T-1}] \\ &= \mathbb{E}[(A_i - \mu)^2 | Y_{i,1}, \dots, Y_{i,T-1}] + \mathbb{E}[U_{i,T+1}^2 | Y_{i,1}, \dots, Y_{i,T-1}] + 2\mathbb{E}[A_i U_{i,T+1} | Y_{i,1}, \dots, Y_{i,T-1}] \\ &= \mathbb{E}[(A_i - \mu)^2 | Y_{i,1}, \dots, Y_{i,T-1}] + \mathbb{E}[U_{i,T+1}^2] + 2\mathbb{E}[A_i \mathbb{E}[U_{i,T+1} | Y_{i,1}, \dots, Y_{i,T-1}, A_i]] \\ &= \mathbb{E}[(A_i - \mu)^2 | Y_{i,1}, \dots, Y_{i,T-1}] + \mathbb{E}[U_{i,T+1}^2] + 2\mathbb{E}[A_i \mathbb{E}[U_{i,T+1}]] \\ &= \mathbb{E}[(A_i - \mu)^2 | Y_{i,1}, \dots, Y_{i,T-1}] + \sigma_i^2.\end{aligned}$$

To obtain the MSFE for IW, first write

$$Y_{i,T+1} - \hat{Y}_{i,T}^{IW} = (Y_{i,T+1} - Y_{i,T}) W_{i,T-1} + (Y_{i,T+1} - \mu)(1 - W_{i,T-1}).$$

Then, we have that

$$\begin{aligned}\text{MSFE}(\text{IW}, \theta_i | Y_{i,1}, \dots, Y_{i,T-1}) &= \mathbb{E} \left[\left(Y_{i,T+1} - \hat{Y}_{i,T}^{IW} \right)^2 | Y_{i,1}, \dots, Y_{i,T-1} \right] \\ &= \mathbb{E} \left[(Y_{i,T+1} - Y_{i,T})^2 | Y_{i,1}, \dots, Y_{i,T-1} \right] W_{i,T-1}^2 + \mathbb{E} \left[(Y_{i,T+1} - \mu)^2 | Y_{i,1}, \dots, Y_{i,T-1} \right] (1 - W_{i,T-1})^2 \\ &\quad + 2\mathbb{E} \left[(Y_{i,T+1} - Y_{i,T})(Y_{i,T+1} - \mu) W_{i,T-1} (1 - W_{i,T-1}) \right] \\ &= 2\sigma_i^2 W_{i,T-1}^2 + \left\{ \mathbb{E}[(A_i - \mu)^2 | Y_{i,1}, \dots, Y_{i,T-1}] + \sigma_i^2 \right\} (1 - W_{i,T-1})^2 \\ &\quad + 2\mathbb{E}[(U_{i,T+1} - U_{i,T})(A_i + U_{i,T+1}) | Y_{i,1}, \dots, Y_{i,T-1}] W_{i,T-1} (1 - W_{i,T-1}) \\ &= 2\sigma_i^2 W_{i,T-1}^2 + \left\{ \mathbb{E}[(A_i - \mu)^2 | Y_{i,1}, \dots, Y_{i,T-1}] + \sigma_i^2 \right\} (1 - W_{i,T-1})^2 + 2\sigma_i^2 W_{i,T-1} (1 - W_{i,T-1}) \\ &= \sigma_i^2 (1 + W_{i,T-1}^2) + \mathbb{E}[(A_i - \mu)^2 | Y_{i,1}, \dots, Y_{i,T-1}] (1 - W_{i,T-1})^2,\end{aligned}$$

which proves the lemma. \square

Proof of Theorem 5.1. To minimize maximum regret, we set

$$\tilde{W}_{i,T-1}^2 = \tilde{W}_{i,T-1} + \tilde{\zeta}_{i,T-1}^2 \left(1 - \tilde{W}_{i,T-1} \right)^2 - \frac{\tilde{\zeta}_{i,T-1}^2}{\tilde{\zeta}_{i,T-1}^2 + 1},$$

equivalently,

$$\tilde{W}_{i,T-1} = 1 - \frac{1}{\sqrt{\tilde{\zeta}_{i,T-1}^2 + 1}}.$$

To verify that \tilde{W}_i is the solution, consider the case that

$$W_{i,T-1} > 1 - \frac{1}{\sqrt{\tilde{\zeta}_{i,T-1}^2 + 1}}.$$

Then,

$$(1 - W_{i,T-1})^2 < \frac{1}{\tilde{\zeta}_{i,T-1}^2 + 1},$$

which in turns implies that

$$\tilde{\zeta}_{i,T-1}^2 (1 - W_{i,T-1})^2 < \frac{\tilde{\zeta}_{i,T-1}^2}{\tilde{\zeta}_{i,T-1}^2 + 1}.$$

Thus, maximum regret is $\sigma_i^2 W_{i,T-1}^2$, which is larger than the solution. Now consider the other case that

$$W_{i,T-1} < 1 - \frac{1}{\sqrt{\tilde{\zeta}_{i,T-1}^2 + 1}}.$$

Now maximum regret is

$$W_{i,T-1}^2 + \tilde{\zeta}_{i,T-1}^2 (1 - W_{i,T-1})^2 - \frac{\tilde{\zeta}_{i,T-1}^2}{\tilde{\zeta}_{i,T-1}^2 + 1}.$$

It remains to show that if $W_{i,T-1} < 1 - \frac{1}{\sqrt{\tilde{\zeta}_{i,T-1}^2 + 1}}$,

$$W_{i,T-1}^2 + \tilde{\zeta}_{i,T-1}^2 (1 - W_{i,T-1})^2 - \frac{\tilde{\zeta}_{i,T-1}^2}{\tilde{\zeta}_{i,T-1}^2 + 1} - \left(1 - \frac{1}{\sqrt{\tilde{\zeta}_{i,T-1}^2 + 1}}\right)^2 > 0.$$

The left-hand side of the inequality above is minimized when $W_{i,T-1} = \tilde{\zeta}_{i,T-1}^2 / (\tilde{\zeta}_{i,T-1}^2 + 1)$, that is, $W_{i,T-1} = 1 - \frac{1}{\sqrt{\tilde{\zeta}_{i,T-1}^2 + 1}}$. This is also the unique minimizer and plugging this value into the left-hand side of the inequality above yields 0. Thus, the left-hand side of the inequality above must be strictly positive if $W_{i,T-1} < 1 - \frac{1}{\sqrt{\tilde{\zeta}_{i,T-1}^2 + 1}}$. Hence, we have proved the desired result. \square

B Feasible Weights for the General Case

This appendix derives the feasible weights in sections 3.2.1 and 3.2.2 in the general case that $TS = \bar{Y}_{i,T} = \Sigma_{t=1}^T Y_{i,t}/T$, $Pool = \mu$ and $IW = \bar{Y}_{i,T} W_{i,T} + \mu(1 - W_{i,T})$.

B.1 IW-O

It is easy to show that the oracle weights that minimize the unconditional MSFE of IW are

$$W_{i,T}^o = \frac{\lambda_i^2}{\lambda_i^2 + \sigma_i^2/T}. \quad (36)$$

These weights follow from equation (9) in Chapter 4 of Timmermann [2006], using the fact that the joint distribution of $Y_{i,T+1}$ and $\hat{Y}_{i,T}^{TS}$ is

$$\begin{pmatrix} Y_{i,T+1} \\ \hat{Y}_{i,T}^{TS} \end{pmatrix} \sim \left(\begin{pmatrix} \mu \\ \mu \end{pmatrix}, \begin{pmatrix} \lambda_i^2 + \sigma_i^2 & \lambda_i^2 \\ \lambda_i^2 & \lambda_i^2 + \sigma_i^2/T \end{pmatrix} \right),$$

which gives the optimal weight on $\hat{Y}_{i,T}^{TS}$ as the product between the inverse of the variance of the forecast and the covariance between the outcome and the forecast, i.e.: $W_{i,T}^o = \frac{\lambda_i^2}{\lambda_i^2 + \sigma_i^2/T}$. The linear combination with $W_{i,T}^o$ as weight could also be obtained by applying the “best linear rule” in equation (9.4), page 129 of Efron and Morris [1973], to $\hat{Y}_{i,T}^{TS}|A_i \sim (A_i, \sigma_i^2/T)$ and $A_i \sim (\mu, \lambda_i^2)$. Feasible oracle weights can for example be obtained as

$$\frac{\Sigma_{t=1}^T (Y_{i,t} - \mu)^2/T - \Sigma_{t=1}^{T-1} (Y_{i,t} - Y_{i,t+1})^2/2(T-1)}{\Sigma_{t=1}^T (Y_{i,t} - \mu)^2/T - \Sigma_{t=1}^{T-1} (Y_{i,t} - Y_{i,t+1})^2/2T}, \quad (37)$$

using the facts that: $\Sigma_{t=1}^T (Y_{i,t} - \mu)^2/T$ is an unbiased estimator of $\lambda_i^2 + \sigma_i^2$; the denominator of the oracle weights can be rewritten as $\lambda_i^2 + \sigma_i^2 - \frac{T-1}{T}\sigma_i^2$ and that $\hat{\sigma}_i^2 = \Sigma_{t=1}^{T-1} (Y_{i,t} - Y_{i,t+1})^2/2(T-1)$ is an unbiased estimator of σ_i^2 . Since both the numerator and the denominator in (37) can be negative, we found that these feasible weights perform very poorly in simulations. We thus considered a number of alternatives and found that the best performance in simulations is obtained by taking the positive part of the numerator in (37) and then again the positive part of the resulting weights, delivering the following feasible weights:²⁶

$$W_{i,T}^{IW-O} = \left(\frac{(\Sigma_{t=1}^T (Y_{i,t} - \mu)^2/T - \Sigma_{t=1}^{T-1} (Y_{i,t} - Y_{i,t+1})^2/2(T-1))^+}{\Sigma_{t=1}^T (Y_{i,t} - \mu)^2/T - \Sigma_{t=1}^{T-1} (Y_{i,t} - Y_{i,t+1})^2/2T} \right)^+. \quad (38)$$

²⁶ Alternatives such as just taking the positive part of the weights in (37) or using the sample covariance between $Y_{i,t}$ and $Y_{i,t-1}$ as an estimator of λ_i^2 in the optimal weights numerator delivered very large errors in simulations.

B.2 IW-MR

To derive the IW-MR weights, we first derive the expression for the conditional MSFEs of TS, Pool and IW.

Lemma B.1. *Let Assumption 3.1 hold. The MSFEs conditional on the information set at time T , $\mathcal{Y}_{N,T}$, are*

$$\begin{aligned}\text{MSFE}(\text{TS}, \theta_i | \mathcal{Y}_{N,T}) &= \sigma_{i,T}^2 + \gamma_{i,T}^2, \\ \text{MSFE}(\text{Pool}, \theta_i | \mathcal{Y}_{N,T}) &= \sigma_{i,T}^2 + \kappa_{i,T}^2, \\ \text{MSFE}(\text{IW}, \theta_i | \mathcal{Y}_{N,T}) &= \sigma_{i,T}^2 + \gamma_{i,T}^2 W_{i,T}^2 + \kappa_{i,T}^2 (1 - W_{i,T})^2 - 2\delta_{i,T} W_{i,T} (1 - W_{i,T}),\end{aligned}$$

where $\sigma_{i,T}^2 = \mathbb{E}[U_{i,T+1}^2 | \mathcal{Y}_{N,T}]$, $\gamma_{i,T}^2 = \mathbb{E}[\bar{U}_{i,T}^2 | \mathcal{Y}_{N,T}]$, $\kappa_{i,T}^2 = \mathbb{E}[(A_i - \mu)^2 | \mathcal{Y}_{N,T}]$ and $\delta_{i,T} = \mathbb{E}[(A_i - \mu)\bar{U}_{i,T} | \mathcal{Y}_{N,T}]$, with $\bar{U}_{i,T} = T^{-1} \sum_{t=1}^T U_{i,t}$.

Proof of Lemma B.1. The MSFE for the TS forecast is given by

$$\begin{aligned}\mathbb{E}[(Y_{i,T+1} - \bar{Y}_{i,T})^2 | \mathcal{Y}_{N,T}] &= \mathbb{E}[(U_{i,T+1} - \bar{U}_{i,T})^2 | \mathcal{Y}_{N,T}] \\ &= \mathbb{E}[U_{i,T+1}^2 | \mathcal{Y}_{N,T}] + \mathbb{E}[\bar{U}_{i,T}^2 | \mathcal{Y}_{N,T}] - 2\mathbb{E}[\bar{U}_{i,T} \mathbb{E}[U_{i,T+1} | \mathcal{Y}_{N,T}, \bar{U}_{i,T}] | \mathcal{Y}_{N,T}] \\ &= \sigma_{i,T}^2 + \gamma_{i,T}^2,\end{aligned}$$

where the last equality follows from

$$\mathbb{E}[U_{i,T+1} | \mathcal{Y}_{N,T}, \bar{U}_{i,T}] = \mathbb{E}[U_{i,T+1} | \mathcal{Y}_{N,T}, A_i] = 0. \quad (39)$$

Now consider the MSFE for the Pool forecast. Note that

$$\begin{aligned}\mathbb{E}[(Y_{i,T+1} - \mu)^2 | \mathcal{Y}_{N,T}] &= \mathbb{E}[(A_i - \mu + U_{i,T+1})^2 | \mathcal{Y}_{N,T}] \\ &= \mathbb{E}[(A_i - \mu)^2 | \mathcal{Y}_{N,T}] + \mathbb{E}[U_{i,T+1}^2 | \mathcal{Y}_{N,T}] + 2\mathbb{E}[(A_i - \mu)U_{i,T+1} | \mathcal{Y}_{N,T}] \\ &= \mathbb{E}[(A_i - \mu)^2 | \mathcal{Y}_{N,T}] + \mathbb{E}[U_{i,T+1}^2 | \mathcal{Y}_{N,T}] + 2\mathbb{E}[A_i \mathbb{E}[U_{i,T+1} | \mathcal{Y}_{N,T}, A_i] | \mathcal{Y}_{N,T}] \\ &= \kappa_{i,T}^2 + \sigma_{i,T}^2,\end{aligned}$$

where the last equality again follows from (39). To obtain the MSFE for IW, first write

$$Y_{i,T+1} - \hat{Y}_{i,T}^{\text{IW}} = (Y_{i,T+1} - \bar{Y}_{i,T}) W_{i,T} + (Y_{i,T+1} - \mu)(1 - W_{i,T}).$$

Then, we have that

$$\begin{aligned}
& \text{MSFE}(\text{IW}, \theta_i | \mathcal{Y}_{N,T}) \\
&= \mathbb{E} \left[\left(Y_{i,T+1} - \widehat{Y}_{i,T}^{\text{IW}} \right)^2 | \mathcal{Y}_{N,T} \right] \\
&= \mathbb{E} \left[\left(Y_{i,T+1} - \bar{Y}_{i,T} \right)^2 | \mathcal{Y}_{N,T} \right] W_{i,T}^2 + \mathbb{E} \left[(Y_{i,T+1} - \mu)^2 | \mathcal{Y}_{N,T} \right] (1 - W_{i,T})^2 \\
&\quad + 2\mathbb{E} \left[(Y_{i,T+1} - \bar{Y}_{i,T}) (Y_{i,T+1} - \mu) W_{i,T} (1 - W_{i,T}) | \mathcal{Y}_{N,T} \right] \\
&= \left\{ \sigma_{i,T}^2 + \mathbb{E} [\bar{U}_{i,T}^2 | \mathcal{Y}_{N,T}] \right\} W_{i,T}^2 + \left\{ \mathbb{E} [(A_i - \mu)^2 | \mathcal{Y}_{N,T}] + \sigma_{i,T}^2 \right\} (1 - W_{i,T})^2 \\
&\quad + 2\mathbb{E} \left[(U_{i,T+1} - \bar{U}_{i,T}) (A_i - \mu + U_{i,T+1}) | \mathcal{Y}_{N,T} \right] W_{i,T} (1 - W_{i,T}) \\
&= [\sigma_{i,T}^2 + \gamma_{i,T}^2] W_{i,T}^2 + [\kappa_{i,T}^2 + \sigma_{i,T}^2] (1 - W_{i,T})^2 \\
&\quad + 2 \left\{ \sigma_{i,T}^2 - \mathbb{E} [(A_i - \mu) \bar{U}_{i,T} | \mathcal{Y}_{N,T}] \right\} W_{i,T} (1 - W_{i,T}) \\
&= \sigma_{i,T}^2 + \gamma_{i,T}^2 W_{i,T}^2 + \kappa_{i,T}^2 (1 - W_{i,T})^2 - 2\delta_{i,T} W_{i,T} (1 - W_{i,T}),
\end{aligned}$$

which proves the lemma. \square

The optimal weights that minimize $\text{MSFE}(\text{IW}, \theta_i | \mathcal{Y}_{N,T})$ are

$$W_{i,T}^* = \frac{\kappa_{i,T}^2 + \delta_{i,T}}{\kappa_{i,T}^2 + \gamma_{i,T}^2 + 2\delta_{i,T}}. \quad (40)$$

We henceforth set $\delta_{i,T} \approx 0$ (this can be seen as approximating $\mathbb{E} [(A_i - \mu) \bar{U}_{i,T} | \mathcal{Y}_{N,T}]$ with the unconditional mean $\mathbb{E} [(A_i - \mu) \bar{U}_{i,T}] = 0$).

Define regret as the difference between the conditional MSFE for a generic weight $W_{i,T}$ and the conditional MSFE that corresponds to the optimal weights $W_{i,T}^*$ in (40):

$$\begin{aligned}
R^*(W_{i,T}, \theta_i | \mathcal{Y}_{N,T}) &:= \text{MSFE}(W_{i,T}, \theta_i | \mathcal{Y}_{N,T}) - \text{MSFE}(W_{i,T}^*, \theta_i | \mathcal{Y}_{N,T}) \\
&= \gamma_{i,T}^2 W_{i,T}^2 + \kappa_{i,T}^2 (1 - W_{i,T})^2 - \frac{\gamma_{i,T}^2 \kappa_{i,T}^2}{\kappa_{i,T}^2 + \gamma_{i,T}^2} \\
&= \gamma_{i,T}^2 \left[W_{i,T}^2 + \zeta_{i,T}^2 (1 - W_{i,T})^2 - \frac{\zeta_{i,T}^2}{\zeta_{i,T}^2 + 1} \right],
\end{aligned} \quad (41)$$

where

$$\zeta_{i,T}^2 := \frac{\kappa_{i,T}^2}{\gamma_{i,T}^2} = \frac{\mathbb{E} [(A_i - \mu)^2 | \mathcal{Y}_{N,T}]}{\mathbb{E} [\bar{U}_{i,T}^2 | \mathcal{Y}_{N,T}]}. \quad (42)$$

The following theorem obtains the optimal minimax regret weight under the assumption that we can put bounds on the random variable $\zeta_{i,T}^2$.

Theorem B.1. *Suppose that it is known that $\zeta_{i,T}^2$ in (42) is such that $\zeta_{i,T}^2 \in [0, \tilde{\zeta}_{i,T}^2]$, where $\tilde{\zeta}_{i,T}^2$*

is large enough that maximum regret can occur at $\zeta_{i,T}^2 = \tilde{\zeta}_{i,T}^2$. Consider maximum regret

$$\max_{\theta_i \in \Theta} R^*(W_{i,T}, \theta_i | \mathcal{Y}_{N,T}) = \gamma_{i,T}^2 \max \left[W_{i,T}^2, \left\{ W_{i,T}^2 + \tilde{\zeta}_{i,T}^2 (1 - W_{i,T})^2 - \frac{\tilde{\zeta}_{i,T}^2}{\tilde{\zeta}_{i,T}^2 + 1} \right\} \right],$$

with $R^*(W_{i,T}, \theta_i | \mathcal{Y}_{N,T})$ defined as in (41). Then, the weight that minimizes maximum regret is

$$\tilde{W}_{i,T} = 1 - \frac{1}{\sqrt{\tilde{\zeta}_{i,T}^2 + 1}}. \quad (43)$$

Proof of Theorem B.1. Same as proof of Theorem 5.1 (with subscript T instead of $T-1$). \square

In applications, the value of the upper bound $\tilde{\zeta}_{i,T}^2$ is uncertain. We thus propose the following heuristic rule to obtain feasible minimax regret optimal weights for IW:

$$\widehat{\tilde{\zeta}_{i,T}^2} := \frac{\max\{(Y_{i,1} - \mu)^2, \dots, (Y_{i,T} - \mu)^2\}}{\sum_{t=1}^{T-1} (Y_{i,t} - Y_{i,t+1})^2 / 2T(T-1)}. \quad (44)$$

Here, the denominator, $\sum_{t=1}^{T-1} (Y_{i,t} - Y_{i,t+1})^2 / 2T(T-1)$, is an unbiased estimator of σ_i^2/T , which approximates $\gamma_{i,T}^2 = \mathbb{E} [\bar{U}_{i,T}^2 | \mathcal{Y}_{N,T}]$ with the unconditional mean $\mathbb{E} [\bar{U}_{i,T}^2] = \sigma_i^2/T$. The numerator, $\max\{(Y_{i,1} - \mu)^2, \dots, (Y_{i,T} - \mu)^2\}$, is a proxy for the upper bound on $\kappa_{i,T}^2 = \mathbb{E} [(A_i - \mu)^2 | \mathcal{Y}_{N,T}]$.

References

- Joshua D. Angrist, Peter D. Hull, Parag A. Pathak, and Cristopher R. Walters. Leveraging lotteries for school value-added: Testing and estimation. *Quarterly Journal of Economics*, 132:871–919, 2017.
- Manuel Arellano and Stephen Bond. Some tests of specification for panel data: Monte carlo evidence and an application to employment equations. *The Review of Economic Studies*, 58(2):277–297, 1991.
- Susan Athey, Raj Chetty, Guido W Imbens, and Hyunseung Kang. The surrogate index: Combining short-term proxies to estimate long-term treatment effects more rapidly and precisely. Working Paper 26463, National Bureau of Economic Research, 2019.
- John M Bates and Clive WJ Granger. The combination of forecasts. *Journal of the Operational Research Society*, 20(4):451–468, 1969.
- Martin Browning, Mette Ejrnaes, and Javier Alvarez. Modelling income processes with lots of heterogeneity. *The Review of Economic Studies*, 77(4):1353–1381, 2010.

- Gary Chamberlain and Keisuke Hirano. Predictive distributions based on longitudinal earnings data. *Annales d'Economie et de Statistique*, pages 211–242, 1999.
- Raj Chetty and Nathaniel Hendren. The impacts of neighborhoods on intergenerational mobility ii: County-level estimates. *The Quarterly Journal of Economics*, 133(3):1163–1228, 2018.
- Raj Chetty, John N. Friedman, and Jonah E. Rockoff. Measuring the impacts of teachers i: Evaluating bias in teacher value-added estimates. *American Economic Review*, 104(9):2593–2632, 2014a. doi: 10.1257/aer.104.9.2593. URL <https://www.aeaweb.org/articles?id=10.1257/aer.104.9.2593>.
- Raj Chetty, John N. Friedman, and Jonah E. Rockoff. Measuring the impacts of teachers ii: Teacher value-added and student outcomes in adulthood. *American Economic Review*, 104(9):2633–79, 2014b. doi: 10.1257/aer.104.9.2633. URL <https://www.aeaweb.org/articles?id=10.1257/aer.104.9.2633>.
- Timothy Christensen, Hyungsik Roger Moon, and Frank Schorfheide. Robust forecasting. arXiv Working Paper arXiv:2011.03153 [econ.EM], 2020. URL <https://arxiv.org/abs/2011.03153>. <https://arxiv.org/abs/2011.03153>.
- Aurore Delaigle. Nonparametric kernel methods with errors-in-variables: constructing estimators, computing them, and avoiding common mistakes. *Australian and New Zealand Journal of Statistics*, 56(2):105–124, 2014.
- Aurore Delaigle and Irene Gijbels. Frequent problems in calculating integrals and optimizing objective functions: a case study in density deconvolution. *Statistics and Computing*, 17: 349–355, 2007.
- Aurore Delaigle, Peter Hall, and Alexander Meister. On deconvolution with repeated measurements. *The Annals of Statistics*, 36(2):665–685, 2008.
- Jeff Dominitz and Charles F. Manski. Minimax-regret sample design in anticipation of missing data, with application to panel data. *Journal of Econometrics*, 2021. doi: <https://doi.org/10.1016/j.jeconom.2020.12.006>. URL <https://www.sciencedirect.com/science/article/pii/S0304407620304000>.
- Bradley Efron. Large-scale inference, Volume 1 of Institute of Mathematical Statistics (ims) Monographs, 2010.
- Bradley Efron. Empirical bayes deconvolution estimates. *Biometrika*, 103(1):1–20, 2016.
- Bradley Efron and Carl Morris. Limiting the risk of Bayes and empirical Bayes estimators—part i: the Bayes case. *Journal of the American Statistical Association*, 66(336):807–815, 1971.

- Bradley Efron and Carl Morris. Stein’s estimation rule and its competitors—an empirical bayes approach. *Journal of the American Statistical Association*, 68(341):117–130, 1973.
- Janqing Fan. On the optimal rates of convergence for nonparametric deconvolution problems. *The Annals of Statistics*, 19:1257-1272, 1991.
- Janqing Fan and Young K. Truong. Nonparametric regression with errors in variables. *The Annals of Statistics*, 21:1900–1925, 1993.
- Jason M. Fletcher, Leora I. Horwitz, and Elizabeth Bradley. Estimating the value added of attending physicians on patient outcomes, National Bureau of Economic Research Working Paper 20534. 2014.
- Jorge Luis García, James J. Heckman, Duncan Ermini Leaf, and María José Prados. Quantifying the life-cycle benefits of an influential early-childhood program. *Journal of Political Economy*, 128(7):2502–2541, 2020.
- Jiaying Gu and Roger Koenker. Unobserved heterogeneity in income dynamics: An empirical bayes perspective. *Journal of Business & Economic Statistics*, 35(1):1–16, 2017. doi: 10.1080/07350015.2015.1052457. URL <https://doi.org/10.1080/07350015.2015.1052457>.
- Fatih Guvenen, Fatih Karahan, Serdar Ozkan, and Jae Song. What do data on millions of us workers reveal about life-cycle earnings risk? Technical report, National Bureau of Economic Research, 2015.
- Peter Hall and Alexander Meister. A ridge-parameter approach to deconvolution. *The Annals of Statistics*, 35(4), 1535-1558, 2007.
- Peter D. Hull. Estimating hospital quality with quasi-experimental data, Working Paper, Brown University. 2020.
- W. James and Charles Stein. Estimation with quadratic loss. proc. 4th Berkeley sympos. math. statist. and prob., vol. i. Berkeley, Calif.: Univ. California press. , 361–379, 1961.
- Thomas J Kane and Douglas O Staiger. Estimating teacher impacts on student achievement: An experimental evaluation. Working Paper 14607, National Bureau of Economic Research, December 2008. URL <http://www.nber.org/papers/w14607>.
- Yong-Seong Kim, Tecla Loup, Joseph Lupton, and Frank P Stafford. Notes on the ‘income plus’ files: 1994-1997 family income and components files. *Documentation, the Panel Study of Income Dynamics* (<http://www.isr.umich.edu/src/psid/income94-97/y-pls-notes.htm>), 2000.
- Patrick Kline, Evan K Rose, and Christopher R Walters. Systemic discrimination among large US employers. *The Quarterly Journal of Economics*, 137(4):1963–2036, 2022.

- Soonwoo Kwon. Optimal shrinkage estimation of fixed effects in linear panel data models, Working Paper. 2021.
- Laura Liu, Hyungsik Roger Moon, and Frank Schorfheide. Forecasting with dynamic panel data models. *Econometrica*, 88(1):171–201, 2020.
- Colin L Mallows and John W Tukey. An overview of techniques of data analysis, emphasizing its exploratory aspects. *Some recent advances in statistics*, 33:111–172, 1982.
- Charles F Manski. Econometrics for decision making: Building foundations sketched by Haavelmo and Wald. *Econometrica*, 89(6):2827–2853, 2021.
- Costas Meghir and Luigi Pistaferri. Income variance dynamics and heterogeneity. *Econometrica*, 72(1):1–32, 2004.
- M Hashem Pesaran, Andreas Pick, and Allan Timmermann. Forecasting with panel data: estimation uncertainty versus parameter heterogeneity, Working Paper 2022.
- James H Stock and Mark W Watson. A comparison of linear and nonlinear univariate models for forecasting macroeconomic time series, In: Engle R, White H Cointegration, Causality and Forecasting: A Festschrift for Clive W.J. Granger. Oxford: Oxford University Press; 1-44, 1999.
- Allan Timmermann. Forecast combinations. *Handbook of Economic Forecasting*, In Elliott, G., C.W.J. Granger and A. Timmermann (Eds.), Handbook of Economic Forecasting. Amsterdam: North Holland, 135-196, 2006