

Heterogeneity in the Persistence of Health: Evidence from a Monthly Micro Panel*

Stephen Hoskins¹, David W. Johnston², Johannes S. Kunz², Michael A. Shields², Kevin E. Staub³

¹Singapore Management University

²Monash University

³University of Melbourne

May 3, 2024

Abstract

Despite being widely used in health economics, dynamic models of health and healthcare typically assume that the persistence in these outcomes is the same for every individual. Understanding the extent and drivers of heterogeneity in persistence is essential for the design and evaluation of health interventions because persistence determines the dynamics and overall long-run effects of such interventions. This paper explores individual-level heterogeneity in the persistence of health outcomes. Using simple regression methods that do not place any restriction on the distribution of the heterogeneity in persistence, the paper documents substantial heterogeneity in health, medical expenditures, and healthcare use. We show that neglecting this heterogeneity leads to estimates that overstate the average persistence and can bias the coefficients of covariates. We find that between 75% and 87% of individuals display persistence and that this persistence is related to the individuals' personality and socio-economic characteristics.

Keywords: self-reported health, healthcare utilisation, dynamic panel data, state dependence, heterogeneity.

JEL classification: I10, I12, C23

**Acknowledgments:* We thank Denzil Fiebig, Tim Halliday, Annika Herr, Andrew Jones, Timothy Neal, Arndt Reichert, Rainer Winkelmann, and participants of the Asian Workshop on Econometrics and Health Economics (Manila), Workshop on the Economics of Health and Wellbeing (Hepburn Springs), Australasian Workshop on Econometrics and Health Economics (Sydney) as well as seminar participants at Hannover, Monash, UTS, and Zurich for valuable feedback. We thank the Center for Research on Successful Ageing (ROSA) at the Singapore Management University for the use of the Singapore Life Panel data. We thank The Ngee Ann Kongsi and the Ministry of Education, Singapore (under its Academic Research Fund Tier 3 program, award reference numbers MOE2013-T3-1-009 and MOE2019-T3-1-006) for financial support.

Authors' addresses: DW Johnston, JS Kunz, MA Shields: Centre for Health Economics, Monash Business School, 900 Dandenong Road, Caulfield, 3145 Victoria, Australia. E-mails: michael.shields@monash.edu, david.johnston@monash.edu, johannes.kunz@monash.edu. KE Staub: Department of Economics, 111 Barry Street, The University of Melbourne, 3010 Victoria, Australia. E-mail: kevin.staub@unimelb.edu.au. S Hoskins: Center for Research on Successful Ageing, Singapore Management University, 10 Canning Rise, #05-01, Singapore 179873

1 Introduction

Health outcome variables, like self-reported health or healthcare expenditures, exhibit a high degree of persistence: health in one period correlates strongly with health in future periods. This paper studies the individual-level heterogeneity in the state dependence of health and health-related outcomes using monthly panel data from a large representative survey of older Singaporeans, the Singapore Life Panel (SLP). Disentangling state dependence from other sources of persistence—such as permanent differences in health levels between individuals—, as well as understanding the extent and drivers of heterogeneity in state dependence, is essential for the design and evaluation of health interventions because state dependence determines the dynamics and overall long-run effects of policies that affect current health. Heterogeneity in state dependence means, for instance, that some individuals recover more slowly from health shocks than others or that some individuals benefit from certain policies for longer than others. Thus, taking into account this heterogeneity among members of the population in their responsiveness to health shocks and policies is a significant issue as it has implications for the efficiency and equitability of interventions.

We propose to model the heterogeneity in state dependence within what is perhaps the most standard framework that has been used in the health economics literature for dynamic panel data models: a linear model with a lagged dependent variable and individual fixed effects (cf., e.g., [Bolhaar, Lindeboom and Van Der Klaauw, 2012](#)). But instead of specifying a single constant coefficient on the lagged dependent variable, we specify a separate coefficient for each individual in the spirit of [Pesaran and Smith \(1995\)](#). The individual-specific state dependence parameters are then estimated freely without imposing any distributional assumptions, akin to the fixed effects estimation of the individual-specific constants. Because the estimator of the state-dependence coefficients, while consistent as the number of time periods increases, is biased in samples with finite time periods, we also explore the performance of bias-corrected estimators that adjust the estimates for this bias.

Dynamic models of health and healthcare are widely used in health economics, with applications ranging from self-assessed health ([Carro and Traferri, 2014](#); [Contoyannis, Jones and Rice, 2004](#)), illness and disability ([Einav, Finkelstein and Mahoney, 2018](#); [Kapteyn, Smith and Van Soest, 2008](#); [Hernandez-Quevedo, Jones and Rice, 2008](#)), to health expenditures ([Kools and Knoef, 2019](#); [Rettenmaier and Wang, 2006](#)), health care use and insurance ([Buchmueller et al., 2021](#); [Bolhaar, Lindeboom and Van Der Klaauw, 2012](#); [Khwaja, 2010](#)), and other health-related behaviours ([Terris-Prestholt and Windmeijer, 2016](#)). Despite differing widely in their applications and methodological approaches, all these papers find evidence for significant state dependence. Yet all of them model state dependence as driven by a single coefficient.¹ To the best of our knowledge only two papers consider heterogeneity

¹Some of the models are nonlinear, in which case the implied marginal effect of a shock to the lagged dependent variable would differ for individuals with different covariate values. However, this heterogeneity is highly restricted and mechanically implied by the model. See the excellent discussion on this issue in [Browning and Carro \(2007\)](#).

in the persistence of health, [Halliday \(2008\)](#) and [Amengual, Bueren and Crego \(2021\)](#). Both pursue approaches that are markedly different from ours and which are based on modelling the distribution of state dependence as consisting of a small number of latent groups. An advantage of our approach is that it makes no assumptions about the shape or nature of the state dependence distribution and can therefore estimate this distribution at a finely granular level. Moreover, the linear panel data model in which it is embedded makes it easy to include a large set of covariates and control for unobserved individual-specific effects, which is not straightforward in the existing approaches.² As the proposed approach requires estimating one state dependence parameter per individual, it is important that the longitudinal dimension of the data is sufficiently large, as we illustrate with simulations. The monthly SLP data in our estimation sample, which has four to eight times more time periods than those used in the mentioned health literature that estimates state dependence heterogeneity, satisfies this requirement comfortably. Indeed, the simulation evidence suggests that some of the yearly household survey panels widely used for health economics research, such as PSID (US Panel Study of Income Dynamics), BHPS (British Household Panel Survey) or SOEP (German Socioeconomic Panel), for which several decades of data exist, would also satisfy it, making our approach potentially broadly applicable.

There are several reasons for modelling state dependence as potentially varying between individuals. One main object of interest in the dynamic health literature cited above is disentangling state dependence from unobserved individual-specific heterogeneity (fixed effects). In practice, this often takes the form of analysing the size and statistical significance of the single coefficient on the lagged dependent variable in a fixed effects panel data model. However, if there actually is heterogeneity in state dependence, this estimator of the autoregressive coefficient is biased for the mean of the distribution, making it impossible to interpret it. Importantly, this is true even if the heterogeneity is completely random and unrelated to any of the regressors in the model or the unobserved individual-specific effects. A second object of interest in the dynamic health literature is to “control” for persistence by ways of the individual-specific effects and the lagged dependent variable in order to estimate the contemporaneous effect of some variable of interest. However, in this case, too, if heterogeneity in state dependence is present in the data generating process, misspecifying the state dependence as constant across individuals leads to bias in the estimator for the coefficient of all other regressors, in general.

Beyond making it possible to analyse the average state dependence and estimate the coefficients of covariates in the case of heterogeneous state dependence, our approach also makes it possible to study further objects of interest which go beyond those of the literature modelling state dependence as constant. In particular, in this paper, the distribution of the state dependence is in itself the key

²For instance, the application in [Halliday \(2008\)](#) contains no controls beyond age and a quadratic function of survey year, and [Amengual, Bueren and Crego \(2021\)](#) does not control for individual-specific effects.

object of interest and can be analysed further, for instance by using the estimated state dependence parameters as outcome variables in regression models.

Methodologically, the approach we propose in this paper is a special case of the mean group estimator introduced in the seminal paper by [Pesaran and Smith \(1995\)](#), which is widely used in macro panels. Perhaps because of past computational and data constraints, the mean group estimator has not been applied to micro panels in health economics: on the one hand, the cross-sectional dimension in such panels is very large (e.g., around ten thousand in our case, while in macro applications it is often around a hundred), while, on the other hand, the longitudinal dimension has in the past been quite short (e.g., five periods in the classic dynamic panel data paper by [Arellano and Bond, 1991](#); while in our case, we use over fifty waves of the SLP). Together, these limitations that arise in the micro panel context mean that many parameters would have to be estimated from very few observations. Here, advances in computation³ and the monthly feature of the SLP data go a long way toward overcoming these obstacles.

However, we also implement two new techniques to account for the dynamic bias, the finite sample bias induced by the lagged dependent variable. First, we apply Empirical Bayes shrinkage ([Chandra et al., 2016](#)) to the estimated state dependence parameters; and second, we use a split-panel jackknife estimator ([Dhaene and Jochmans, 2015](#)) to obtain bias-corrected estimates of the state dependence parameters as in [Chudik and Pesaran \(2019\)](#). Empirical Bayes shrinkage is a popular technique used in the applied microeconomic literature to deal with the imprecision of the estimates of individual-specific constants when analysing these effects is of primary interest (such as hospital or doctor fixed effects). Here, we extend this approach to the shrinkage of heterogeneous slopes. [Chudik and Pesaran \(2019\)](#) provide theoretical results and simulation evidence on the split-panel approach to a heterogeneous dynamic panel data model such as the one in this paper. We extend the simulation evidence to study data generating processes with covariates (and therefore to be able to consider the finite sample properties of the estimators of the coefficients on covariates) and apply this method for the first time to an empirical application. Our simulation evidence suggests that with the number of time periods available in our application, the dynamic bias is not a first-order concern, with the mean group procedure delivering satisfactory estimates. However, the Empirical Bayes and the bias-correction approach improve the estimator's performance further, and they could be critical for applications with fewer time periods.

There is an active theoretical literature on heterogeneous panels. Two papers closely related to ours are [Mavroeidis, Sasaki and Welch \(2015\)](#) and [Okui and Yanagi \(2019\)](#). The first considers a model similar to ours, but in a setting where the longitudinal dimension is short and fixed. To estimate the heterogeneity of state dependence consistently in this setting, they develop a nonparametric approach.

³For instance, sets of fixed effects in linear models can be partialled-out and estimated efficiently with moderate computing power (e.g. [Correia, 2016](#)).

In contrast, in our setting the longitudinal dimension is long and we can rely on the substantially simpler methods of [Pesaran and Smith \(1995\)](#) and [Chudik and Pesaran \(2019\)](#). The second paper considers the case without covariates and is therefore essentially model-free. Their estimator is also equal to the mean group estimator of [Pesaran and Smith \(1995\)](#) on which they apply a modified [Dhaene and Jochmans \(2015\)](#) bias correction. Our approach is similar, but we operate in a model-based framework with covariates entering the model through a linear index.

We apply the methods to the monthly SLP survey from Singapore, which includes detailed health, health care and health expenditure questions in over 50 waves from roughly 9,000 individuals. As mentioned, this higher frequency elicitation of various health behaviours and outcomes makes it possible to assess persistence and adaptation patterns in greater detail than in previous studies. We compare standard dynamic panel data methods (such as random effects, naïve fixed effects and jackknife-corrected fixed effects, which all constrain the persistence to a single parameter) with models that allow for fully-flexible heterogeneous response parameters across individuals. Estimates of population-average persistence are obtained as simple means as well as means based on Empirical Bayes shrinkage.

We have five main findings. First, there is substantial heterogeneity across individuals in the state dependence of self-assessed health, healthcare use (as measured by at least one doctor visit) and out-of-pocket healthcare expenditures. Second, this heterogeneity drastically biases standard random-effects and fixed-effect estimators, including those that correct for the dynamic bias. The biases can be as large as 100%. Third, the heterogeneity in state dependence is only weakly correlated across the health and healthcare outcomes we considered, but in each case there is strong evidence of correlation with baseline characteristics. Fourth, personality traits have more predictive power among these baseline characteristics than economic status (income) or economic preferences (risk). Fifth, in all three outcomes, we find that about 13% to 25% of individuals display little to no state dependence, with the remaining majority of individuals having a roughly bell-shaped distribution with the majority of its mass in the positive support region.

The remainder of this paper is structured as follows. [Section 2](#) discusses the econometric methods by introducing panel data models with heterogeneous dynamics and considering possible estimation approaches ([Section 2.1](#)), and investigating the finite sample performance of the estimators via simulation experiments ([Section 2.2](#)). [Section 3](#) contains our analysis of the estimated heterogeneity in the persistence of health and healthcare: we give details on the Singapore Life Panel survey and the estimation sample ([Section 3.1](#)) and present the results of our estimations ([Section 3.2](#)). [Section 4](#) concludes.

2 Econometric methods

2.1 Fixed effects estimation of panel data models with heterogenous dynamics

Consider the following stylised dynamic panel model for individual i 's health or healthcare utilisation in period t :

$$y_{it} = \alpha_i + \rho y_{it-1} + x'_{it}\beta + \varepsilon_{it}, \quad i = 1, \dots, N, \quad t = 1, \dots, T, \quad (1)$$

where α_i are individual-specific effects, x_{it} is a vector of K covariates, and ε_{it} is an IID idiosyncratic error term that is uncorrelated with x_{it} and α_i . The α_i are treated as fixed effects in that no distributional assumption is placed on them; in particular, they may be correlated with the covariates. The use of this fixed effects auto-regressive (AR) panel model is widespread in health economics to model dynamics in health or healthcare use, which are captured by the parameter ρ . The main estimation problem with (1) discussed in the literature has been the so-called Nickell bias or dynamic bias (Nickell, 1981) that arises when the total number of time periods T is small and treated as fixed. To see the problem, consider for instance that after applying the within transformation to equation (1) to rid it of the α_i the demeaned y_{it-1} is correlated with the demeaned ε_{it} . The default approach has been to use IV and GMM approaches in which (functions of) earlier lags of the dependent variable (and covariates) are used as instruments (Anderson and Hsiao, 1981; Arellano and Bond, 1991; Arellano and Bover, 1995; Blundell and Bond, 1998, among others). Since such procedures can suffer from the many/weak instruments problem, alternatively one can relax the view of T as fixed. Asymptotic expansions show that the dynamic bias, which is of order T , can be removed in a number of ways; for instance, by using the split-panel jackknife (SPJ) estimator (Dhaene and Jochmans, 2015; Chudik, Pesaran and Yang, 2018).

Dynamic panel data models such as (1) have been used to separate different sources of persistence in health, such as “individual heterogeneity” (as captured by the individual effects α_i) and state dependence (as captured by the lagged dependent variable). The main object of interest of this paper is to examine the individual heterogeneity *in* state dependence, that is, the distribution of state dependence over individuals. To that end, we consider a heterogeneous version of (1), which we will refer to as the heterogeneous auto-regressive (HAR) panel model, in which state dependence is individual-specific, too:

$$y_{it} = \alpha_i + \rho_i y_{it-1} + x'_{it}\beta + \varepsilon_{it}, \quad (2)$$

where in contrast to (1) this model explicitly specifies individual heterogeneity in the coefficient of the lagged dependent variable (ρ_i). The default in the dynamic health literature is to assume that

$\rho = \rho_i$ for all individuals, ignoring this type of heterogeneity. While in this paper the goal is to study the heterogeneity in itself and its correlates with observables, sometimes the interest lies only in the *average* state dependence, $\bar{\rho} = N^{-1} \sum_{i=1}^N \rho_i$. In that case, often the estimates $\hat{\rho}$ from homogenous AR models such as (1) are treated as estimates of $\bar{\rho}$. However, it is known that $\hat{\rho}$ from model (1) is not consistent for $\bar{\rho}$ (Pesaran and Smith, 1995), and, as we illustrate in simulations below, even in very simple data generating processes the estimates of models that ignore heterogeneity in state dependence can be severely biased. Thus, if one wrongly assumes $\rho_i = \rho$ and estimates model (1), there are now two sources of bias: the Nickell bias and the heterogeneity bias.

We propose to implement and compare three solutions to this problem. First, since the dynamic bias dissipates as T increases, but not the heterogeneity bias, a simple approach is to neglect the dynamic bias and focus only on the heterogeneity bias by estimating each separate ρ_i and then averaging over all N estimates. This is the Mean Group (MG) estimator of Pesaran and Smith (1995), only that the vector β is homogenous. In the context of the SLP data, this approach relies on the substantial longitudinal dimension of about 50 periods, which could imply that the residual dynamic bias might be negligible. To implement this approach, we estimate model (2) by pooled OLS of y_{it} on x_{it} , y_{it-1} , a set of individual indicator variables (to estimate α_i), and a set of individual indicator variables interacted with y_{it-1} (to estimate ρ_i). While estimating a model with in the order of $2N$ parameters is computationally intensive, all parameters of the HAR model can be estimated fast and efficiently using the computational approach of Correia (2016), which relies on a combination of alternating symmetric projections, the FrischWaughLovell theorem and conjugate gradient acceleration. Using the estimates, the average state dependence is obtained as

$$\hat{\rho}^{MG} = \frac{1}{N} \sum_{i=1}^N \hat{\rho}_i.$$

In a second approach, we explore potential improvements in the estimation quality of $\bar{\rho}$ that could be gained from refining the procedure of averaging over the $\hat{\rho}_i$: Instead of the simple average as in the MG estimator, we also obtain an estimate of $\bar{\rho}$ by applying Empirical Bayes (EB) shrinkage to the $\hat{\rho}_i$ when averaging them. EB methods are often used to estimate fixed effects such as α_i . They consist in applying to the individual estimates an adjustment which shrinks them towards the average in a way that is, loosely speaking, proportional to their variance. This adds bias to the individual estimates but reduces their variance, often resulting in improved mean squared error. In our context, we propose to obtain the average state dependence as

$$\hat{\rho}^{EB} = \frac{1}{N} \sum_{i=1}^N \hat{\rho}_i^{EB}, \quad \text{where } \hat{\rho}_i^{EB} = (1 - \hat{w}_i) \hat{\rho}_i + \hat{w}_i \hat{\mu}_\rho,$$

$\hat{\mu}_\rho$ is an estimate of $\mu_\rho = E(\rho_i)$ and the terms \hat{w}_i are estimates of $w_i = \frac{V(\hat{\rho}_i)}{V(\hat{\rho}_i) + \sigma_\rho^2}$, where $V(\hat{\rho}_i)$ is the

variance of $\hat{\rho}_i$, σ_ρ^2 is the variance of ρ_i , and $\mu_\rho = E(\rho_i)$.⁴ While in the standard case of unbiased estimators the EB procedure unambiguously adds bias, here it can also reduce the bias of some individual state dependence parameters, which in turn may lead to a reduction of the bias in the average $\hat{\rho}^{EB}$.⁵

While these two approaches take into account the heterogeneity in ρ_i , they do not correct estimates for possible Nickell and small- T bias. The third approach is also based on the MG approach, but instead of applying the averaging to estimates of a naïve fixed effects estimator, we obtain the estimates for ρ_i from a bias-corrected (BC) estimator that addresses the dynamic bias: the split-panel jackknife estimator for heterogeneous panel models proposed by [Chudik and Pesaran \(2019\)](#). This estimator extends the estimator of [Dhaene and Jochmans \(2015\)](#)—which nonparametrically removes the Nickell bias and other finite sample bias due to T being finite—to heterogeneous models such as HAR. In our empirical analysis, we use first-order split-panel jackknife (also called half-panel jackknife) implementations of these estimators, which consist in splitting the estimation sample into two halves at the middle period $T/2$ and estimating the model in each half, as well as in the full sample. By comparing the average estimate in the half-sized samples to the estimate in the full sample, it is possible to infer the magnitude of the order- T bias (which is larger in the small samples and decreases with T) and subtract the bias from the full sample estimate. For model (2), this means that we obtain the average state dependence as

$$\hat{\rho}^{BC} = \frac{1}{N} \sum_{i=1}^N \hat{\rho}_i^{BC}, \quad \text{where } \hat{\rho}_i^{BC} = 2\hat{\rho}_i - (\hat{\rho}_i^{FH} + \hat{\rho}_i^{SH})/2,$$

and $\hat{\rho}_i^{FH}$ and $\hat{\rho}_i^{SH}$ are estimates of ρ_i obtained from using the first and second half-samples, respectively.

Heterogeneous panel models such as (2) are common in empirical macroeconomics ([Breitung, 2015](#)). While they have been developed for and applied in macroeconomic settings, much less is known about the performance of these models in micro panels. An essential feature which makes these models attractive for health economics is that in order to get purely data-driven and consistent estimates of ρ_i (as well as α_i) none of the heterogeneity is constrained by distributional assumptions. That is, these terms are treated as fixed effects to be estimated. However, different from the default macroeconomic setup, where i often indexes countries, in the context of the Singapore Life Panel the individual dimension N is (very) large. Another difference is that in the macroeconomic approach the interest sometimes lies in long-term changes in the outcome (such as $\rho_i/(1 - \rho_i)$ and $\beta/(1 - \rho_i)$) rather than

⁴The EB-shrunken estimates $\hat{\rho}_i^{EB}$ are estimated by an iterative procedure due to [Morris \(1983\)](#) which alternates between estimating w_i for given (μ_ρ, σ_ρ) , and (μ_ρ, σ_ρ) for given w_i . We use the implementation of this procedure by [Chandra et al. \(2016\)](#).

⁵Since the bias in the estimated ρ_i is towards 0, the EB shrinkage will reduce the bias for estimators of ρ_i that lie between 0 and $\bar{\rho}$. Because the EB shrinkage is proportional to the variance of the estimator, and the variance of $\hat{\rho}_i$ decreases with the absolute value of ρ_i , EB can result in a reduced average bias for a number of distributions of ρ_i , such as e.g. symmetric or right-skewed distributions over the positive support of ρ_i (0,1), or symmetric or left-skewed distributions over the negative support of ρ_i (-1,0).

the short-term dynamics (ρ_i and β). These long-term objects of interest are often estimated directly by transforming the model into what is called its error-correction form and using maximum likelihood estimation to estimate the resulting nonlinear model. In contrast, as mentioned, we estimate model (2) for the Singapore Life Panel directly without transformation using what is arguably the simplest possible approach: OLS with sets of indicator or dummy variables to capture individual (and time) effects, and interactions between a set of individual dummies and the lagged dependent variable to estimate ρ_i .

While the default in the macroeconomic literature has been to ignore the dynamic bias, the results for the dynamic bias in the AR panel model (1) directly extend to the HAR model. As in the AR model, the dynamic bias in the autoregressive coefficients (ρ_i) spills over to the covariate coefficients in the HAR model (2). And as in the AR model, the dynamic bias increases when fewer covariates with $\beta \neq 0$ are included in the HAR model. In contrast, in the health literature, the default has been to ignore the heterogeneity bias. As mentioned, if an AR model is assumed when the DGP is a HAR model, the estimated $\hat{\rho}$ is inconsistent for $E(\rho_i)$. In general, the asymptotic bias in $\hat{\rho}$ spills over to the covariates, biasing $\hat{\beta}$. This will generally happen even when the ρ_i are statistically independent of α_i and x_{it} . Thus, the presence of heterogeneity in state dependence has important implications for the conventional strategy in health economics of studying heterogeneity via interactions (for instance, by estimating an augmented version of model (1) with the additional interaction variable $y_{it-1}d_i$, where d_i is an indicator of some group membership such as gender). Unless the group indicator in the interaction (d_i) fully captures the heterogeneity, the coefficient estimates on the auto-regressive terms and covariates will still be inconsistent. The approach suggested in this paper, by contrast, provides a straightforward way of obtaining consistent estimates of group differences in average state dependence: first, estimate all ρ_i ; then, regress the estimated ρ_i on any regressors of interest such as d_i .

2.2 Numerical experiments

DATA GENERATING PROCESS AND ESTIMATORS

To investigate the finite sample performance of the various estimators for the HAR model (2) that we have discussed above we design a series of Monte Carlo simulation experiments for settings similar to our monthly micro data. We use a simple data generating process (DGP) corresponding to (2) with a single regressor x_{it} . Similar to the Singapore Life Panel data used, the number of periods is set to $T = 50$. We only use $N = 500$, since the quality of individual-specific parameters such as ρ_i is

determined by T . The model is

$$y_{i0} \sim \text{Normal}(0, 0.25) \quad (3)$$

$$y_{it} = \alpha_i + \rho_i y_{it-1} + 0.5x_{it} + \varepsilon_{it}, \quad \text{if } t \geq 1, \quad (4)$$

for $i = 1 \dots, 500$ and $t = 1, \dots, 50$. The error is drawn from $\varepsilon \sim \text{Normal}(0, 0.5)$ and $x_{it} = \bar{x}_i + \nu + 0.001t^2$, where $\nu \sim \text{Normal}(0, 0.25)$. The variables $(\alpha_i, \rho_i, \bar{x}_i)$ follow a trivariate normal distribution with means $(1, 0.5, 1)'$, and α_i and \bar{x}_i have standard deviations $\sigma_\alpha = 0.25$ and $\sigma_{\bar{x}} = 1$. The parameters that we vary in the simulation experiments are the heterogeneity in the auto-regressive parameter ρ_i . We consider different standard deviations, $\sigma_\rho = \{0, 0.125, 0.25\}$ as well as a case in which the variance of ρ_i is heteroskedastic with $E(\sigma_{\rho,i}) = 0.15$. In this case, we draw ρ_i from five latent groups with standard deviations $0.05, 0.10, \dots, 0.25$. Each latent group has the same probability mass. The second set of parameters that we vary is the correlation between α_i, ρ_i , and \bar{x}_i ,

$$C_{\alpha\rho\bar{x}} = \begin{pmatrix} 1, & \tau_{\alpha\rho}, & \tau_{\alpha\bar{x}} \\ \tau_{\alpha\rho}, & 1, & \tau_{\bar{x}\rho} \\ \tau_{\alpha\bar{x}}, & \tau_{\bar{x}\rho}, & 1 \end{pmatrix}, \quad (5)$$

where, for instance, $\tau_{\alpha\rho} = \sigma_{\alpha\rho}/(\sigma_\alpha\sigma_\rho)$ and $\sigma_{\alpha\rho}$ is the covariance between α_i and ρ_i , and the other parameters in (5) are defined analogously. We consider a baseline DGP with $C_{\alpha\rho\bar{x}} = C^B$ where α_i, ρ_i , and \bar{x}_i are uncorrelated and an alternative DGP with $C_{\alpha\rho\bar{x}} = C^A$ with non-zero correlations:

$$C^B = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad C^A = \begin{pmatrix} 1 & -0.3 & -0.8 \\ -0.3 & 1 & 0.5 \\ -0.8 & 0.5 & 1 \end{pmatrix}. \quad (6)$$

Further correlation structures are considered in robustness checks and discussed below.

We compare six possible estimators in the Monte Carlo simulation focusing on the question of how well they perform in obtaining an estimate of $E(\rho_i)$, the mean of the autoregressive coefficient. The first three ignore the individual heterogeneity in ρ_i ; that is, they assume $\rho_i = \rho$ for all i as in equation (1), and we use $\hat{\rho}$ as their estimate of the mean, $E(\rho_i)$. The latter three obtain an estimate of each ρ_i , which is then combined into an average $\bar{\rho}$.

The estimators in the first category which ignores heterogeneity are the Random Effects estimator (RE), the Fixed Effects estimator (FE) and a Bias-Corrected Fixed Effects estimator (BC-FE). For this last estimator we use the first-order split-sample jackknife estimator of [Dhaene and Jochmans \(2015\)](#). The three estimators that take into account the individual heterogeneity in ρ_i are the Mean Group estimator (MG), the Mean Group estimator based on Empirical-Bayes shrinkage of ρ_i (EB-MG), and the Mean Group estimator based on split-sample-jackknife Bias Correction (BC-MG).

RESULTS FOR $E(\rho_i)$

The results of the Monte Carlo experiment are presented in Tables 1, for the baseline DGP where the ρ_i are completely random and uncorrelated to both α_i and the observable covariate, and Table 2, for the alternative DGP where the ρ_i are correlated with the individual fixed effects α_i and the covariate x_{it} . In the top half, the tables' columns report estimates of $E(\rho_i)$ for each of the six different estimators across four rows representing four different DGPs that vary in terms of the heterogeneity in ρ_i as captured by its standard deviation, σ_ρ . In the top row, there is no heterogeneity (“*Panel A*: $\sigma_\rho = 0$ ”) and the model simplifies to the standard constant-effects dynamic panel data model from equation (1) as assumed by the estimators in columns (1)-(3). The lower half of the tables presents the corresponding estimates for β . For both parameters β and $E(\rho)$, the true value is 0.5.

A look at the top row (“*Panel A*: $\sigma_\rho = 0$ ”) of Table 1, shows that the random effects estimator is biased even in this simple DGP ($\hat{\rho} = 0.66$). If ρ were equal to zero, RE would be the most efficient estimator under this DGP. But when $\rho \neq 0$ the lagged dependent variable, which contains α_i , is mechanically correlated with the estimation error, which also contains α_i in the RE approach. All other estimators condition on α_i and so do not have this problem. Their estimates are all close to the true value of 0.5, although the FE estimator and the MG estimator (and therefore the EB-MG estimator, too) suffer from slight attenuation bias due to some remaining Nickell bias. The estimators that correct for this via the split-panel jackknife (BC-FE and BC-MG) are virtually unbiased.

However, the introduction of even a modest level of heterogeneity in ρ_i (*Panel B*) has quite dramatic effects on the estimates of the constant-effects model estimators of columns (1)-(3), with biases ranging from over 60 to close to 100 percent (i.e., $0.812/0.5$ and $0.969/0.5$). Biases for these estimators worsen further for increased levels of heterogeneity (*Panel C*) or a similar level of average heterogeneity but with heteroskedasticity (*Panel D*). These results are quite remarkable, since estimates of ρ close to 1 or even greater than 1 (as in *Panel C*) could lead a researcher not aware of potential heterogeneity in ρ to erroneously believe that the model is not stationary ($\rho > 1$), when in reality every single ρ_i is in fact smaller than 1. In contrast, in each of these cases, the simple Mean Group estimator (column 4) delivers a good performance with a downward bias that is less than 10 percent. The more sophisticated Mean Group estimators EB-MG and BC-MG obtain estimates that are even less biased.

RESULTS FOR β

The lower half of Table 1 contains the corresponding results for the estimates of the coefficient on the covariate, β . The results mirror those for $E(\rho)$. The performance of all estimators except for RE is good in the case of no heterogeneity in ρ (*Panel E*). But with heterogeneity in ρ_i all the estimators that assume no heterogeneity produce strongly distorted estimates of β , even though ρ_i is uncorrelated with both α_i and x_{it} . Similar to the case of the estimates of $E(\rho_i)$ before, the biases can easily be so

large as to lead to seriously misleading inference. For instance, in *Panel G*, the average RE estimate of β is essentially zero (-0.005), when the true value is 0.5. Researchers relying on BC-FE estimates might even conclude that the effect of x_{it} is negative (-0.087). In contrast, the estimators that account for the heterogeneity in ρ_i produce estimates of β that are close to the true value.⁶

ROBUSTNESS CHECKS

To obtain the results in Table 2, the ρ_i were drawn as being correlated with α_i and x_{it} according to the correlation matrix C^A in (6). The table shows that the presence of such correlation does not change any of the results substantially. The estimates are quite similar. We investigated further correlation structures (results not reported) where we changed both the sign and strength of the correlation. In all cases, the results were very similar to the ones reported in Tables 1 and 2. We also considered a DGP with the special case of the (demeaned) x_{it} being independent over time, $x_{it} = \bar{x}_i + \nu$, instead of serially correlated as in our baseline. In this case, the biases are limited to $E(\rho_i)$ (and similar to the ones reported in the tables) but do not spill over to β . Unfortunately, in many real-world applications x_{it} are serially correlated, so heterogeneity in ρ_i might be an issue for researchers even if their interest is limited to β . Finally, we also repeated the simulations for a different, lower value of $E(\rho_i) = 0.2$, which is close to our results in the application. Again, results were qualitatively similar.

SIMULATIONS INVESTIGATING THE ROLE OF T

We conclude the simulations by exploring the role of T in obtaining reliable estimates of the average state dependence parameter, $E(\rho_i)$. Results are collected in Figure 1. We present estimates for the FE, MG and BC-MG estimators for, successively, $T=6, 10, 25, 50$. The smallest possible time period is $T = 6$ as BC-MG halves the sample along the time dimension, resulting in samples of three periods, and one further period is lost due to the lagged dependent variable. The right-hand-side panel of Figure 1 introduces an additional empirical feature: missing observations. In the estimates of that panel, 20 per cent of the observations were deleted at random. In our DGP, small- T bias (which always leads to attenuation bias) and heterogeneity bias (which here leads to overestimation) work in opposing directions. Thus, with a low number of time periods, these biases almost cancel out for FE. However, already at $T=10$, the heterogeneity bias dominates, and FE is severely distorted. In contrast, the BC-MG approach delivers very good results from as little as 10 time periods. The simple mean group estimator without bias-correction requires a larger longitudinal dimension for the dynamic bias to dissipate naturally. As seen before, at $T=50$, the remaining bias is small.

CONCLUSION

To summarise, our simulations have illustrated that with the number of periods at hand, neglecting

⁶There are no results for MG-EB since there is no heterogeneity in β and therefore the MG-EB estimate for β is the same as the one in column (4).

heterogeneity in ρ_i leads to substantial biases, whereas the main concern typically raised in dynamic fixed- T panel data, Nickell bias, play a more limited role. While the more sophisticated approaches can serve to further address these, especially when T is smaller, the gains from using these methods over more simple and robust approaches such as the mean group estimator are more modest when T is around 50, as in our empirical application.

3 Estimates of the heterogeneity in the persistence of health and healthcare

3.1 The Singapore Life Panel: Data and institutional setting

Data for this analysis is drawn from the Singapore Life Panel (SLP), a panel survey which has been collecting data from a population-representative internet survey of older Singaporeans, operating every month since 2015.

Initial recruitment to the SLP took place from May to July 2015 and utilized a random sample of 4% of all households that contained at least one Singaporean citizen or permanent resident aged 50 to 69 years of age. Traditional methods were used to invite these individuals and their spouses to participate in the panel, including introductory letters followed-up by door-knocking and phone calls. Of the 22,500 addresses contacted, individuals from 11,500 households elected to join the panel, implying a 52% response rate. Including spouses, 16,000 respondents joined the panel.

This base sample are highly representative of the corresponding national population for this age group on demographic and socioeconomic characteristics including gender, marital status, ethnicity, labour force participation, income and expenditure (Vaithianathan et al., 2018). However, there has been a mild attrition effect in each month since then, mostly driven by respondents with more education being more likely to complete subsequent surveys.

Every month since September 2015, respondents have been invited to log into the online platform and participate in a 15 minute survey in exchange for SGD\$10 of supermarket vouchers. In a typical month, 7,700 respondents will complete the survey: 69% of their own accord, 25% with help from a family member, 5% with the help over the phone from a Singapore Management University student, and 1% at an in-person appointment. We use data through to December 2019, and restrict the sample to those individuals who have answered at least 10 waves of data and who provided non-missing responses to the key variables below. This combines to roughly 350,000 observations, from 8,500 individuals over 53 months. Due to lower attrition among respondents with high education, this restricted sample is slightly over-representative of both the highly-educated and Chinese populations. For example, 35%

of our sample have a post-secondary qualification, compared to 30% among the equivalent age cohort in the Singapore population. Similarly, 86% in our sample are Chinese, compared to 79% of the wider population [Department of Statistics \(2010\)](#). Age and gender remain highly population-representative.

Questions repeated every month include employment status, sources of income, expenditure across 44 categories, health status and subjective well-being. We make use of three health outcome variables in this paper:

1. Self-assessed bad health. We construct this variable as an average across two related measures: *Self-assessed health*, measured by asking “*Would you say your health is excellent, very good, good, fair, or poor?*,” and *health satisfaction* assessed on a five-point Likert scale from “*very dissatisfied*” to “*very satisfied*.” We rescale the resulting variable, which has 10 points of support, to the unit interval ($y_{it} = 0, 0.1, 0.2, \dots, 1$); where 0 indicates the best and 1 the worst health.⁷

2. Medical expenditures. Measured in Singapore dollars, we use this variable to capture behavioural responses in healthcare utilisation. Medical expenditures are the sum of monthly spending on medications, outpatient services, hospital services and home nursing. It includes both out-of-pocket expenditure and any expenses paid from MediSave, the mandatory individual health savings account.

3. Doctor visits. A binary indicator variable of whether the individual visited a doctor in the past month, mostly free of charge.

Figure 2 presents the overall distributions of each of these three variables (top row) and the distributions of their individual averages (bottom row), visualising the variation in our health outcome measures. One can see from the individual averages that some individuals display no variation in one or more of the outcome variables (that is, when the individual mean is zero in any of these outcomes or, for bad health and doctor visits, one). For such individuals with “perfect” persistence, α_i and ρ_i are not identified and no estimates of their state dependence, $\hat{\rho}_i$, are obtained.

The healthcare use and expenditures take place within a healthcare system which is primarily government-run, but funded via individual mandatory health savings accounts and private health insurance, supplemented by private healthcare providers and targeted support. All citizens and permanent residents must contribute 8% to 10.5% of their wages to a national savings account called MediSave, which is administered by the Central Provident Fund (CPF), a government entity ([Central Provident Fund Board, 2020](#)). Funds from this account must be used to purchase MediShield Life, a basic health insurance plan which covers large medical bills from public hospitals ([Ministry of Health, 2020](#)). Individuals can also use their MediSave funds to purchase approved supplementary private health insurance, and to cover approved health expenditures up to posted withdrawal limits. For those Singaporeans who cannot otherwise afford healthcare, assistance is provided via Medifund and targeted subsidies.

⁷We code having higher values of the variable as indicating worse health in order to align it with the other two outcomes (medical expenditures and doctor visits) where higher values also indicate higher health-related needs.

3.2 Estimation results

To estimate the heterogeneity in the state dependence of each of the three outcome variables (self-assessed bad health, medical expenditures and doctor visits) we estimate for each of them model (1) and (2) by RE, FE, BC-FE, the simple mean group (MG) estimator, as well as the empirical-Bayes and bias-corrected MG estimators (EB-MG and BC-MG). Besides individual-specific constants and the lagged dependent variable, the specification includes year fixed effects, month-of-the-year fixed effects, and a quadratic function of age. Standard errors are clustered at the household level.

ESTIMATES OF THE AVERAGE STATE DEPENDENCE

Table 3 presents the estimates in a format that is similar to the one used for the simulation results, focusing on the estimation of $E(\rho_i)$, the average state dependence in health, expenditures and doctor visits. The pattern of the results in Table 3 is also in line with the simulations in each of the panels *A* to *C* corresponding to the three outcomes. First, the random effects estimator of the homogeneous ρ model (1) (column 1) yields results that are strikingly different from the other estimators, which all condition on individual-specific constants α_i and therefore are robust against correlation between α_i and the regressors. The RE estimates of state dependence are about twice as large as those of the the two remaining estimators of the homogeneous ρ model, FE and BC-FE in columns (2) and (3), and range from 0.859 for self-assessed bad health to 0.299 for medical expenditures. Second, FE and BC-FE produce estimates of the average state dependence that are noticeably larger than those from the estimators that account for heterogeneity in ρ_i , shown in columns (4)-(6). Because the dynamic bias correction increases the estimate, this overestimation of the (average) state dependence is more pronounced for BC-FE than for FE.

Third, the estimators of the heterogeneous state dependence model produce results that are consistent with each other. For *self-assessed bad health*, the simple and Empirical-Bayes-shrunken MG estimate are both 0.237, while the split-panel jackknife bias-corrected estimate of the average state dependence is 0.268. These estimates are thus only about half as large as the BC-FE estimate. While the BC-MG estimator has favourable theoretical properties, a practical disadvantage is that data on individuals with gaps in their responses have to be dropped from the estimation (that is, individuals who skipped their interview in one or more months). This is why the number of observations is lower for BC-MG. This is also the likely reason why there is a slightly larger discrepancy between the MG estimates and the BC-MG estimate for *medical expenditures* in *Panel B*, where the MG and EB-MG estimates are 0.111 and 0.131 while BC-MG is 0.040, although with a standard error that makes it not statistically different to the MG estimates. Again, the BC-FE estimate is about twice that of the other MG estimates. For *doctor visits* in *Panel C*, the three heterogeneous ρ_i estimates are again very close to each other at 0.150, 0.150 and 0.154, while BC-FE is about 50 percent larger. Fourth, the table

reports the share of the estimated ρ_i that are statistically significant. Despite the estimated average state dependence being of a moderate level, most $\hat{\rho}_i$ are statistically significant. This ranges from 74.3 percent for medical expenditure to 94.3 and 96.4 percent for self-assessed bad health and doctor visits, respectively.

In conclusion, estimating the heterogeneous ρ_i model, we find strong evidence of state dependence in each of the outcomes. However, neglecting this heterogeneity and estimating a constant- ρ model would substantially overestimate the average state dependence in these health outcomes. Since we found that the simple, uncorrected MG estimates of the heterogeneous model are largely similar to the bias-corrected ones and that they have the advantage of handling gaps in the observations more easily, we use these estimates as the basis for investigating the distribution of the estimated state dependence further.

THE DISTRIBUTION OF THE ESTIMATED STATE DEPENDENCE

Figure 3 presents histograms of the estimated $\hat{\rho}_i$ for each of the three health outcomes. Despite having been estimated freely and without any constraint, the shapes of the three histograms display broad similarities. While the outcomes capture very different aspects of healthcare, such as subjective health, healthcare cost, and healthcare use, all three histograms are bell-shaped and have a spike in density around the value of 0.⁸ The vast majority of the density mass is positive or zero, with only a small fraction being negative, mainly on the interval $(-0.5, 0)$. While expenditures and visits have also little mass beyond 0.5, the figure shows that self-assessed bad health has a larger mass with such high persistence.

We performed two robustness checks on the estimation of the distribution of the heterogeneity in the state dependence of these health outcomes. First, the heterogeneous auto-regressive model (2) with one single lag of the dependent variable is a stylised representation of the dynamics in these health outcomes. To make sure that by limiting the model to one lag we are not missing crucial higher-order dynamics, we re-estimated the models for each of the three outcomes including additional lags of the dependent variable. The resulting distributions of persistence are similar (cf. Figure A1 in the appendix), reassuring us that the simple HAR model with one lag captures the essential differences in dynamics across individuals. Second, since the quality of the estimates of ρ_i depends on the number of periods available per individual, a concern is that lower-quality estimates of ρ_i from individuals observed for fewer periods distort the shape of the estimated histogram. To address this, we plotted the histograms keeping only the $\hat{\rho}_i$ of individuals observed for 30 or more periods in the data. The resulting distributions of persistence (cf. Figure A2) are also virtually identical to the ones in Figure 3.

⁸As mentioned previously, for individuals with no variation in the outcome, ρ_i is not identified. That is, the distribution shown in the figure is conditional on individuals displaying some time-variation in y_{it} .

To investigate whether individual-level persistence is correlated across the outcome domains, we compute pairwise correlation coefficients, such as $Corr(\hat{\rho}_i^{SAH}, \hat{\rho}_i^{Op})$, the correlation between the estimated state dependence in self-assessed bad health and the estimated state dependence in medical expenditures. Moreover, similarly, we can also determine how individual state dependence ρ_i correlates with the individual constants or “fixed effects” α_i . The latter represent permanent or underlying levels of the health-related outcomes (net of covariates), while the former give an indication of how fast or slow an individual returns to these underlying levels after a shock. Table 4 presents the full correlation matrix between these two heterogeneity terms.

Considering the lower part of Table 4 which contains the correlations between the ρ_i , the results indicate that there is a weak positive correlation between state dependence in self-assessed bad health and state dependence in doctor visits (0.118) as well as between state dependence in doctor visits and state dependence in medical expenditures (0.089). The state dependence in health and expenditures are essentially uncorrelated (0.016). These contrasts with the correlation among individual averages or long-term individual levels, α_i , where we find consistently strong positive correlations. The correlation between the long-term health level and the corresponding long-term expenditure and doctor visit levels are 0.517 and 0.543. The correlation in α_i^{Exp} and α_i^{Doc} is 0.968. Of the cross-correlations between state dependence and permanent levels, most are close to zero, with two exceptions. First, there is a strong negative correlation between state dependence and levels in self-assessed bad health (-0.597). Sadly, this implies that a good shock in self-assessed health for those in long-term bad health tends to be short-lived, while bad shocks for those in long-term good health tend to be persistent. Second, there is a slight negative correlation between state dependence in doctor visits and long-term levels of self-assessed bad health (-0.073): those in good long-term health tend to display a somewhat stronger persistence in doctor visits.

THE SOCIAL GRADIENT IN STATE DEPENDENCE

Is the heterogeneity in state dependence random? If not, can it be predicted by observables? Which traits are associated with high state dependence? To answer these questions we use regression models where the estimated $\hat{\rho}_i$ are the dependent variables. As potential determinants of the state dependence we consider socio-economic variables such as gender (an indicator variable for female), marital status (an indicator for being married), a quadratic function in age, years of schooling and (the natural logarithm of) household income. We also consider risk attitudes as a measure of economic preferences, and variables for the “big five” personality traits.

To take into account the distributions of the estimated state dependence parameters (Figure 3), which are bell-shaped but with a prominent spike around zero, we use finite mixture modelling. We assume that the observed distribution is a mix of two latent (or unobserved) normally-distributed groups or classes. One class produces the overall bell shape with positive mean and most mass in the positive

domain. The other class produces the spike around zero as a normal distribution centred at zero and with a small variance. This modelling is consistent with the idea of the distribution of ρ_i having a discrete mass point at zero. Since we observe estimates of the state dependence ($\hat{\rho}_i$) rather than the true ρ_i , the estimation error in $\hat{\rho}_i$ naturally transforms the true mass point at zero of ρ_i into a distribution around zero for $\hat{\rho}_i$. Finally, each of the two classes has a different conditional mean function specified as a linear index in the covariates mentioned above. Formally, let the probability density function of $\hat{\rho}_i$ be

$$f(\hat{\rho}_i|x_i) = \pi f_1(\hat{\rho}_i|x_i) + (1 - \pi) f_2(\hat{\rho}_i|x_i), \quad (7)$$

where

$$f_1(\hat{\rho}_i|x_i) \sim N(x_i'\gamma_1, \sigma_1) \quad \text{and} \quad f_2(\hat{\rho}_i|x_i) \sim N(x_i'\gamma_2, \sigma_2), \quad (8)$$

and $\pi \in (0,1)$ denotes the proportion of class 1 in the population. We estimate the parameters of model (7)-(8) by maximum likelihood based on the sample log likelihood $\mathcal{L}_n(\pi, \gamma_1, \gamma_2, \sigma_1, \sigma_2) = \sum_{i=1}^n \ln f(\hat{\rho}_i|x_i)$.

We consider three specifications for the two mean functions, $E(\hat{\rho}_i|x_i) = x_i\gamma_1$ (for those i in class 1) and $E(\hat{\rho}_i|x_i) = x_i\gamma_2$ (for those in class 2). In the first specification, we exclude all variables x_i and the conditional means are just constants ($\gamma_{1,0}$ and $\gamma_{2,0}$); in the second, we include socio-economic variables in x_i but exclude the economic preference and personality variables; and in the third, we include both socio-economic variables as well as economic preferences and personality. The results for these three specifications, for each of the three health outcomes, are given in Tables 5-7. In the tables, classes are always ordered such that the class with the higher conditional mean is defined as ‘class 2’.

From the goodness-of-fit criteria reported in the tables—the Akaike Information Criterion (AIC) and the Bayesian Information Criterion (BIC)—it is immediately clear that the full specification with both sets of covariates, socioeconomic and preferences/personality, is the preferred one as it produces the smallest values of the criteria and thus the best fit (while penalising models for larger numbers of parameters). Using likelihood ratio tests, we can also formally test for joint significance of these sets of variables. For instance, both the set of socioeconomic characteristics (gender, marital status, age, education and income) as well as the set of preferences/personality are each highly jointly statistically significant ($p < 0.001$) for predicting the persistence in self-assessed health (columns 1 and 2) and in medical expenditures (columns 3 and 4). For doctor visits (columns 5 and 6), we find no statistically significant effect of socioeconomic characteristics, but again the set of preferences/personality is highly significant.

Figure 4 presents a further visual confirmation of the goodness-of-fit of the finite mixture model for self-assessed bad health from Table 7. Each ρ_i has been predicted as either class 1 (if $\hat{\pi}_i > 0.5$) or class 2 (if $\hat{\pi}_i \leq 0.5$) and the resulting histograms of the two predicted classes are presented on the

left-hand-side graph. From the graph it is clear that the class assignment coincides with the intuitive and assumed division of classes into the spike at zero (class 1) and the bell-shape over mostly the positive domain of ρ_i (class 2). The right-hand-side graph plots the point-wise marginal predicted density (that is, $\hat{f}(\hat{\rho}_i|x)$ from equation 7), which indicates that the estimated model also replicates well the shape of the raw empirical distribution of the data. The corresponding figures for the other two outcomes look similar and can be found in the appendix (cf. Figure A3).

Thus, overall, the findings show that the heterogeneity in the state dependence in health, expenditures, and doctor visits is not random. Socioeconomic variables as well as personality traits are statistically significant predictors of this heterogeneity and improve the fit of the estimation; in the case of personality variables, substantially so. At the level of individual predictors the picture is less precise, with several variables not being individually statistically significant. Of course, for class 1, we expect the variables to have coefficients equal to zero since we hypothesise $E(\hat{\rho}_i) = 0$, consistent with the spike at zero. Indeed, the coefficients in columns (1), (3) and (5) of Table 7 are almost all insignificant, and the class means are close to zero (-0.027, -0.024 and -0.054).

For the coefficients of the mean function of class 2 (i.e. for individuals with non-zero persistence), the most precise results are obtained for health expenditures (column 4), where higher age, education and income are all associated with a higher average state dependence. This might be an indication that when individuals with a higher socio-economic status are hit with a health shock, they use more or longer follow-up treatments and tests associated with out-of-pocket costs. Apart from this case, however, economic variables such as household income and risk attitudes are not found to be individually significant in the preferred models (Table 7). This contrasts with the personality traits, where individually significant variables are found in every outcome (columns 2, 4 and 6).

4 Conclusions

Methodologically, the focus of much of the research using dynamic health economics models has been on addressing Nickell bias. As both longer and higher-frequency micro panel datasets are increasingly becoming available to health economists, we argue that there should be a shift in the attention of researchers towards also addressing heterogeneity bias, an issue which so far has been largely neglected in the dynamic health literature. As we demonstrated with a monthly panel survey from Singapore, there is extensive evidence of heterogeneity in state dependence in common and widely-used outcome variables in this literature, such as subjective health, medical expenditures and doctor visits. In our empirical application, for instance, we found that neglecting heterogeneity led to overestimating the average state dependence between 50% to 100%.

Heterogeneity in state dependence is different from heterogeneity in the effect of control variables in

that if neglected it generally leads to biases in both the estimate of the average state dependence as well as the coefficients of control variables even if the heterogeneity is fully independent of the covariates or the individual-specific effects. A consequence of this is that heterogeneity in state dependence should not be investigated in the usual way by creating interactions between the dependent variable and variables of interest. Instead, we propose to freely estimate the heterogeneity first (that is, without imposing any distributional assumptions on the heterogeneity), and then, in a second step, analyse how the heterogeneity relates to observables of interest. Using this approach with the Singaporean data, we found that the state dependence in health outcomes was correlated with socioeconomic characteristics, personality traits and economic preferences.

The proposed estimation methods are easy to implement and are now computationally feasible even in large datasets. They are based on methods that are well-established in econometrics and macroeconomics (Pesaran and Smith, 1995), and which have been further refined recently (Chudik and Pesaran, 2019). This paper applied these methods for the first time to dynamic health models. Our simulations showed that with the data at hand the basic Mean Group estimator performs well and is effective at estimating heterogeneity. For shorter data, adding bias-correction and empirical-Bayes-shrinkage adjustments to Mean Group estimation can substantially improve the performance. While in macroeconomics the focus has been mainly on the average state dependence and, in particular, the average long-run effect, in our context we have focused on further analysing the heterogeneity. We found that between 13% and 25% of individuals had no state dependence in the health outcomes we considered, while the remainder’s distribution of state dependence was mainly positive and bell-shaped. The results showed state dependence to be only weakly correlated across health outcomes and revealed different patterns of correlation between the heterogeneity in state dependence and the heterogeneity in long-term levels depending on the outcome. Such associations of state dependence with socioeconomic and other characteristics, as well as its correlation with long-term levels, are important from a policy perspective since they have implications for the efficiency and equitability of interventions targeting these health outcomes by amplifying or reducing differences between population subgroups.

References

- Amengual, Dante, Jesús Bueren and Julio A Crego. 2021. “Endogenous health groups and heterogeneous dynamics of the elderly.” *Journal of Applied Econometrics* 36(7):878–897.
- Anderson, Theodore Wilbur and Cheng Hsiao. 1981. “Estimation of dynamic models with error components.” *Journal of the American Statistical Association* 76(375):598–606.
- Arellano, Manuel and Olympia Bover. 1995. “Another look at the instrumental variable estimation of error-components models.” *Journal of Econometrics* 68(1):29–51.

- Arellano, Manuel and Stephen Bond. 1991. "Some tests of specification for panel data: Monte Carlo evidence and an application to employment equations." *Review of Economic Studies* 58(2):277–297.
- Blundell, Richard and Stephen Bond. 1998. "Initial conditions and moment restrictions in dynamic panel data models." *Journal of Econometrics* 87(1):115–143.
- Bolhaar, Jonneke, Maarten Lindeboom and Bas Van Der Klaauw. 2012. "A dynamic analysis of the demand for health insurance and health care." *European Economic Review* 56(4):669–690.
- Breitung, Joerg. 2015. The Analysis of Macroeconomic Panel Data. In *The Oxford Handbook of Panel Data*. Oxford University Press.
<https://doi.org/10.1093/oxfordhb/9780199940042.013.0015>
- Browning, Martin and Jesus Carro. 2007. *Heterogeneity and Microeconometrics Modeling*. Vol. 3 of *Econometric Society Monographs* Cambridge University Press p. 4774.
- Buchmueller, Thomas C, Terence C Cheng, Ngoc TA Pham and Kevin E Staub. 2021. "The effect of income-based mandates on the demand for private hospital insurance and its dynamics." *Journal of Health Economics* 75:102403.
- Carro, Jesús M and Alejandra Traferri. 2014. "State dependence and heterogeneity in health using a bias-corrected fixed-effects estimator." *Journal of Applied Econometrics* 29(2):181–207.
- Central Provident Fund Board. 2020. "MediSave."
<https://www.cpf.gov.sg/Members/Schemes/schemes/healthcare/medisave>
- Chandra, Amitabh, Amy Finkelstein, Adam Sacarny and Chad Syverson. 2016. "Health Care Exceptionalism? Performance and Allocation in the US Health Care Sector." *American Economic Review* 106(8):2110–2144.
- Chudik, Alexander and M Hashem Pesaran. 2019. "Mean group estimation in presence of weakly cross-correlated estimators." *Economics Letters* 175:101–105.
- Chudik, Alexander, M Hashem Pesaran and Jui-Chung Yang. 2018. "Half-panel jackknife fixed-effects estimation of linear panels with weakly exogenous regressors." *Journal of Applied Econometrics* 33(6):816–836.
- Contoyannis, Paul, Andrew M Jones and Nigel Rice. 2004. "The dynamics of health in the British Household Panel Survey." *Journal of Applied Econometrics* 19(4):473–503.
- Correia, Sergio. 2016. "A feasible estimator for linear models with multi-way fixed effects." *Working paper, available from <http://scorreia.com/research/hdfe.pdf>*.
- Department of Statistics. 2010. Census of Population. Statistical Release 1 Demographic Characteristics, Education, Language and Religion. Technical report Ministry of Trade & Industry, Republic

of Singapore.

https://www.singstat.gov.sg/-/media/files/publications/cop2010/census_2010_release1/cop2010sr1.pdf

- Dhaene, Geert and Koen Jochmans. 2015. “Split-panel jackknife estimation of fixed-effect models.” *Review of Economic Studies* 82(3):991–1030.
- Einav, Liran, Amy Finkelstein and Neale Mahoney. 2018. “Provider incentives and healthcare costs: Evidence from long-term care hospitals.” *Econometrica* 86(6):2161–2219.
- Halliday, Timothy J. 2008. “Heterogeneity, state dependence and health.” *Econometrics Journal* 11(3):499–516.
- Hernandez-Quevedo, Cristina, Andrew M Jones and Nigel Rice. 2008. “Persistence in health limitations: a European comparative analysis.” *Journal of Health Economics* 27(6):1472–1488.
- Kapteyn, Arie, James P Smith and Arthur Van Soest. 2008. “Dynamics of work disability and pain.” *Journal of Health Economics* 27(2):496–509.
- Khwaja, Ahmed. 2010. “Estimating willingness to pay for medicare using a dynamic life-cycle model of demand for health insurance.” *Journal of Econometrics* 156(1):130–147.
- Kools, Lieke and Marike Knoef. 2019. “Health and consumption preferences; estimating the health state dependence of utility using equivalence scales.” *European Economic Review* 113:46–62.
- Mavroeidis, Sophocles, Yuya Sasaki and Ivo Welch. 2015. “Estimation of heterogeneous autoregressive parameters with short panel data.” *Journal of Econometrics* 188(1):219–235.
- Ministry of Health. 2020. “MediShield Life.”
<https://www.moh.gov.sg/cost-financing/healthcare-schemes-subsidies/medishield-life>
- Morris, Carl N. 1983. “Parametric empirical Bayes inference: theory and applications.” *Journal of the American Statistical Association* 78(381):47–55.
- Nickell, Stephen. 1981. “Biases in dynamic models with fixed effects.” *Econometrica* 49(6):1417–1426.
- Okui, Ryo and Takahide Yanagi. 2019. “Panel data analysis with heterogeneous dynamics.” *Journal of Econometrics* 212(2):451–475.
- Pesaran, M Hashem and Ron Smith. 1995. “Estimating long-run relationships from dynamic heterogeneous panels.” *Journal of Econometrics* 68(1):79–113.
- Rettenmaier, Andrew J and Zijun Wang. 2006. “Persistence in Medicare reimbursements and personal medical accounts.” *Journal of Health Economics* 25(1):39–57.

Terris-Prestholt, Fern and Frank Windmeijer. 2016. “How to sell a condom? The impact of demand creation tools on male and female condom sales in resource limited settings.” *Journal of Health Economics* 48:107–120.

Vaithianathan, Rhema, Bryce Hool, Michael D Hurd and Susann Rohwedder. 2018. “High-frequency internet survey of a probability sample of older Singaporeans: the Singapore Life Panel®.” *Singapore Economic Review* p. 1842004.

Tables and figures

Table 1: MONTE CARLO SIMULATION RESULTS: ZERO CORRELATION

I. Average estimate of $E(\rho_i)$. (True value in DGP: $E(\rho_i) = 0.500$)						
	Estimators for AR models			Estimators for HAR models		
	RE	FE	BC-FE	MG	EB-MG	BC-MG
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: $\sigma_\rho = 0$</i>						
y_{it-1}	0.663	0.477	0.503	0.457	0.471	0.500
(sd)	(0.002)	(0.001)	(0.001)	(0.001)	(0.000)	(0.001)
<i>Panel B: $\sigma_\rho = 0.125$</i>						
y_{it-1}	0.964	0.812	0.969	0.456	0.472	0.498
(sd)	(0.008)	(0.022)	(0.030)	(0.002)	(0.003)	(0.003)
<i>Panel C: $\sigma_\rho = 0.250$</i>						
y_{it-1}	1.027	0.983	1.038	0.454	0.469	0.492
(sd)	(0.001)	(0.002)	(0.004)	(0.002)	(0.002)	(0.003)
<i>Panel D: $E(\sigma_{\rho,i}) = 0.150$</i>						
y_{it-1}	1.017	0.951	1.048	0.459	0.474	0.496
(sd)	(0.002)	(0.006)	(0.004)	(0.000)	(0.000)	(0.001)
II. Average estimate of β . (True value in DGP: $\beta = 0.500$)						
	Estimators for AR models			Estimators for HAR models		
	RE	FE	BC-FE	MG	EB-MB	BC-MG
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel E: $\sigma_\rho = 0$</i>						
x_{it}	0.344	0.521	0.509	0.529		0.523
(sd)	(0.002)	(0.004)	(0.002)	(0.005)		(0.003)
<i>Panel F: $\sigma_\rho = 0.125$</i>						
x_{it}	0.066	0.245	0.059	0.529		0.520
(sd)	(0.008)	(0.021)	(0.030)	(0.005)		(0.003)
<i>Panel G: $\sigma_\rho = 0.250$</i>						
x_{it}	-0.005	0.097	-0.087	0.524		0.514
(sd)	(0.001)	(0.003)	(0.002)	(0.005)		(0.002)
<i>Panel H: $E(\sigma_{\rho,i}) = 0.150$</i>						
x_{it}	0.010	0.122	-0.074	0.535		0.529
(sd)	(0.003)	(0.006)	(0.005)	(0.006)		(0.004)

Note: Cells contain average and standard deviation of estimates over 100 replications. In each replication, observations are drawn from the following DGP: $y_{i,t} = \alpha_i + \rho_i y_{i,t-1} + x'_{i,t} \beta + \varepsilon_{i,t}$, $i = 1, \dots, 500$, $t = 1, \dots, 50$, $\alpha_i \sim N(1, 0.25^2)$, $\rho_i \sim N(0.5, \sigma_\rho^2)$, $\beta = 0.5$, $x_{it} = \bar{x}_i + \nu$, $\bar{x}_i \sim N(0, 1)$, $\nu \sim N(0, 0.025^2)$, $\varepsilon_{i,t} \sim N(0, 0.5^2)$, $Corr(\alpha_i, \rho_i) = -0.3$, $Corr(\rho_i, \bar{x}_i) = 0.5$, $Corr(\alpha_i, \bar{x}_i) = -0.8$, $y_{i0} \sim N(0, 0.5^2)$. In the rare instance that a particular draw of ρ_i is realised outside of the interval $(-0.99, 0.99)$, it is recoded as either -0.99 or 0.99, whichever is closer.

Table 2: MONTE CARLO SIMULATION RESULTS: NON-ZERO CORRELATION

I. Average estimate of $E(\rho_i)$. (True value in DGP: $E(\rho_i) = 0.500$)						
	Estimators for AR models			Estimators for HAR models		
	RE	FE	BC-FE	MG	EB-MG	BC-MG
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: $\sigma_\rho = 0$</i>						
y_{it-1}	0.615	0.475	0.500	0.455	0.469	0.499
(sd)	(0.001)	(0.001)	(0.001)	(0.001)	(0.000)	(0.001)
<i>Panel B: $\sigma_\rho = 0.125$</i>						
y_{it-1}	0.994	0.889	1.033	0.461	0.477	0.505
(sd)	(0.011)	(0.026)	(0.032)	(0.001)	(0.001)	(0.002)
<i>Panel C: $\sigma_\rho = 0.250$</i>						
y_{it-1}	1.027	0.989	1.028	0.460	0.475	0.498
(sd)	(0.000)	(0.002)	(0.002)	(0.003)	(0.003)	(0.003)
<i>Panel D: $E(\sigma_{\rho,i}) = 0.150$</i>						
y_{it-1}	1.019	0.961	1.042	0.459	0.474	0.498
(sd)	(0.002)	(0.008)	(0.003)	(0.003)	(0.004)	(0.004)
II. Average estimate of β . (True value in DGP: $\beta = 0.500$)						
	Estimators for AR models			Estimators for HAR models		
	RE	FE	BC-FE	MG	EB-MG	BC-MG
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel E: $\sigma_\rho = 0$</i>						
x_{it}	0.295	0.527	0.505	0.536		0.520
(sd)	(0.001)	(0.001)	(0.001)	(0.001)		(0.001)
<i>Panel F: $\sigma_\rho = 0.125$</i>						
x_{it}	0.052	0.169	-0.042	0.532		0.517
(sd)	(0.010)	(0.025)	(0.036)	(0.002)		(0.002)
<i>Panel G: $\sigma_\rho = 0.250$</i>						
x_{it}	0.015	0.079	-0.095	0.530		0.511
(sd)	(0.001)	(0.002)	(0.001)	(0.002)		(0.002)
<i>Panel H: $E(\sigma_{\rho,i}) = 0.150$</i>						
x_{it}	0.018	0.102	-0.087	0.531		0.519
(sd)	(0.002)	(0.007)	(0.003)	(0.007)		(0.009)

Note: Cells contain average and standard deviation of estimates over 100 replications. In each replication, observations are drawn from the following DGP: $y_{i,t} = \alpha_i + \rho_i y_{i,t-1} + x'_{i,t} \beta + \varepsilon_{i,t}$, $i = 1, \dots, 500$, $t = 1, \dots, 50$, $\alpha_i \sim N(1, 0.25^2)$, $\rho_i \sim N(0.5, \sigma_\rho^2)$, $\beta = 0.5$, $x_{it} = \bar{x}_i + \nu$, $\bar{x}_i \sim N(0, 1)$, $\nu \sim N(0, 0.025^2)$, $\varepsilon_{i,t} \sim N(0, 0.5^2)$, $Corr(\alpha_i, \rho_i) = -0.3$, $Corr(\rho_i, \bar{x}_i) = 0.5$, $Corr(\alpha_i, \bar{x}_i) = -0.8$, $y_{i0} \sim N(0, 0.5^2)$. In the rare instance that a particular draw of ρ_i is realised outside of the interval $(-0.99, 0.99)$, it is recoded as either -0.99 or 0.99, whichever is closer.

Table 3: ESTIMATION RESULTS OF DIFFERENT PANEL ESTIMATORS

<i>Dependent variables: Various measures of health and health behaviours</i>						
	Estimators for AR models			Estimators for HAR models		
	RE	FE	BC-FE	MG	EB-MG	BC-MG
	(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A: Self-assessed bad health [0-1]</i>						
y_{it-1}	0.859 (0.003)	0.345 (0.005)	0.468 (0.005)	0.237 (0.003)	0.237 (0.003)	0.268 (0.007)
Obs Estimation sample	330,636	330,636	330,636	330,636	330,636	316,182
Mean y	0.483	0.483	0.483	0.483	0.483	0.483
Obs Grouped mean				8,200	8,199	5,664
Share significant				0.943		
<i>Panel B: Medical expenditures [0,inf)</i>						
y_{it-1}	0.299 (0.020)	0.182 (0.019)	0.217 (0.019)	0.111 (0.008)	0.131 (0.006)	0.040 (0.073)
Obs Estimation sample	318,467	318,467	318,467	318,467	318,467	304,883
Mean y	143.541	143.541	143.541	146.230	146.230	149.576
Obs Grouped mean				7,827	7,826	5,634
Share significant				0.743		
<i>Panel C: Doctor visit [0,1]</i>						
y_{it-1}	0.413 (0.004)	0.168 (0.003)	0.223 (0.003)	0.150 (0.003)	0.150 (0.003)	0.154 (0.005)
Obs Estimation sample	316,619	316,619	316,619	316,619	316,619	303,069
Mean y	0.343	0.343	0.343	0.348	0.348	0.348
Obs Grouped mean				7,688	7,687	4,942
Share significant				0.964		
Year fixed effects	✓	✓	✓	✓	✓	✓
Month fixed effects	✓	✓	✓	✓	✓	✓
Quadratic in age	✓	✓	✓	✓	✓	✓
Clustering household-level	✓	✓	✓	✓	✓	✓

Notes: Table shows estimation results of the various linear dynamic panel data models shown in (1) and (2), its coefficients and clustered standard errors on the household level in brackets. Columns display various dynamic estimation approaches, (1) Random Effects, (2) Fixed Effects, (3) Bias-corrected FE. Cols. (4), (5), (6) estimate a $\hat{\rho}_i$ for each person and average them: (4) Mean Group, (5) Empirical-Bayes-adjusted MG, (6) Bias-corrected MG.

Source: SLP v53, own calculations.

Table 4: CORRELATION MATRIX BETWEEN ESTIMATES OF INDIVIDUAL STATE DEPENDENCE (ρ_i) AND INDIVIDUAL LEVELS (α_i) ACROSS HEALTH OUTCOMES

Variables	α_i^{SAH}	α_i^{Exp}	α_i^{Doc}	ρ_i^{SAH}	ρ_i^{Exp}	ρ_i^{Doc}
α_i^{SAH}	1					
α_i^{Exp}	0.517 (0.000) [8,603]	1				
α_i^{Doc}	0.543 (0.000) [8,610]	0.968 (0.000) [8,603]	1			
ρ_i^{SAH}	-0.597 (0.000) [8,200]	0.020 (0.069) [8,195]	-0.009 (0.436) [8,180]	1		
ρ_i^{Exp}	0.003 (0.795) [7,827]	-0.035 (0.002) [7,827]	-0.017 (0.137) [7,816]	0.016 (0.166) [7,495]		
ρ_i^{Doc}	-0.073 (0.000) [7,684]	-0.016 (0.173) [7,681]	0.006 (0.603) [7,666]	0.118 (0.000) [7,367]	0.089 (0.000) [7,263]	1

Notes: Cells contain: in boldface, pairwise correlation coefficients for α_i and ρ_i ; in parentheses, p-values for the test of correlation being equal to 0; and in brackets, the number of observations used for the test. Outcome abbreviations: SAH – self-assessed bad health, Exp – medical expenditures, Doc – doctor visits.

Source: SLP v53, own calculations.

Table 5: ECONOMIC GRADIENTS IN STATE DEPENDENCE, FINITE MIXTURE MODEL, CONSTANTS-ONLY SPECIFICATION

<i>Dependent variables: Estimated individual state dependence $\hat{\rho}_i$</i>						
	Subjective bad health		Medical expenditures		Doctor visits	
	Class 1	Class 2	Class 1	Class 2	Class 1	Class 2
	(1)	(2)	(3)	(4)	(5)	(6)
N	8,153		7,683		7,671	
AIC	1,964.205		-118.752		-1,518.802	
BIC	1,999.236		-84.019		-710.167	
Class probabilities	0.812	0.188	0.228	0.772	0.165	0.835
Class means	0.155	0.581	-0.026	0.142	-0.056	0.189
Class σ	0.053	0.028	0.002	0.078	0.002	0.060

Notes: Estimates from model (7)-(8). AIC: Akaike Information Criterion. BIC: Bayesian Information Criterion. Class probabilities: Estimates of π (class 1) and $1 - \pi$ (class 2).

Source: SLP v53, own calculations.

Table 6: ECONOMIC GRADIENTS IN STATE DEPENDENCE, FINITE MIXTURE MODEL, SOCIO-ECONOMIC VARIABLES SPECIFICATION

<i>Dependent variables: Estimated individual state dependence $\hat{\rho}_i$</i>						
	Subjective bad health		Medical expenditures		Doctor visits	
	Class 1	Class 2	Class 1	Class 2	Class 1	Class 2
	(1)	(2)	(3)	(4)	(5)	(6)
Female (Yes/No)	-0.002 (0.005)	-0.008 (0.007)	-0.001 (0.004)	-0.000 (0.008)	0.000 (0.004)	0.008 (0.007)
Married (Yes/No)	-0.002 (0.006)	-0.008 (0.009)	-0.001 (0.006)	-0.008 (0.010)	-0.006 (0.007)	0.010 (0.009)
Age in years/10	-0.066 (0.106)	0.104 (0.150)	-0.080 (0.086)	0.418 (0.167)	0.063 (0.100)	0.150 (0.146)
Age squared	0.005 (0.009)	-0.009 (0.013)	0.008 (0.007)	-0.034 (0.014)	-0.005 (0.008)	-0.014 (0.012)
Years of schooling/10	-0.005 (0.007)	0.020 (0.009)	0.014 (0.005)	0.029 (0.010)	0.003 (0.005)	-0.018 (0.009)
log HH income	-0.006 (0.002)	0.002 (0.004)	0.003 (0.002)	0.013 (0.004)	0.002 (0.002)	0.001 (0.004)
N	7,818		7,379		7,357	
AIC	1,498.263		-152.162		-700.083	
BIC	1,616.655		-34.753		-582.725	
Class probabilities	0.111	0.889	0.228	0.772	0.165	0.835
Class means	-0.028	0.268	-0.027	0.142	-0.055	0.189
Class σ	0.002	0.074	0.002	0.077	0.002	0.060

Notes: Estimates from model (7)-(8). Standard errors robust to heteroskedasticity and clustering at the household level. AIC: Akaike Information Criterion. BIC: Bayesian Information Criterion. Class probabilities: Estimates of π (class 1) and $1 - \pi$ (class 2).

Source: SLP v53, own calculations.

Table 7: ECONOMIC GRADIENTS IN STATE DEPENDENCE, FINITE MIXTURE MODEL, FULL SPECIFICATION

<i>Dependent variables: Estimated individual state dependence $\hat{\rho}_i$</i>						
	Subjective bad health		Medical expenditures		Doctor visits	
	Class 1	Class 2	Class 1	Class 2	Class 1	Class 2
	(1)	(2)	(3)	(4)	(5)	(6)
Female (Yes/No)	-0.002 (0.007)	-0.014 (0.008)	-0.002 (0.005)	0.003 (0.009)	-0.004 (0.005)	0.009 (0.008)
Married (Yes/No)	-0.000 (0.007)	-0.013 (0.010)	-0.003 (0.006)	-0.016 (0.011)	-0.008 (0.008)	-0.001 (0.010)
Age in years/10	0.061 (0.120)	0.192 (0.163)	-0.013 (0.093)	0.315 (0.185)	0.072 (0.109)	0.050 (0.153)
Age squared	-0.006 (0.010)	-0.016 (0.014)	0.002 (0.008)	-0.025 (0.016)	-0.006 (0.009)	-0.005 (0.013)
Years of schooling/10	-0.003 (0.009)	0.015 (0.010)	0.012 (0.006)	0.021 (0.012)	0.001 (0.005)	-0.003 (0.010)
log HH income	-0.003 (0.003)	0.002 (0.004)	0.001 (0.002)	0.016 (0.005)	0.001 (0.002)	0.004 (0.004)
Risk attitudes (1-10)	0.008 (0.015)	0.008 (0.016)	0.001 (0.009)	-0.001 (0.018)	-0.002 (0.009)	0.012 (0.015)
Extraversion (std)	0.011 (0.006)	0.005 (0.008)	-0.004 (0.004)	-0.011 (0.009)	-0.007 (0.005)	0.022 (0.007)
Agreeableness (std)	0.010 (0.008)	0.006 (0.008)	0.003 (0.004)	0.020 (0.009)	-0.002 (0.006)	-0.016 (0.008)
Conscientiousness (std)	-0.008 (0.008)	0.018 (0.008)	0.008 (0.004)	0.003 (0.009)	0.003 (0.007)	-0.017 (0.008)
Neuroticism (std)	0.007 (0.007)	0.030 (0.007)	0.003 (0.004)	0.023 (0.008)	-0.002 (0.005)	0.004 (0.007)
Openness (std)	-0.010 (0.008)	-0.011 (0.008)	-0.002 (0.005)	0.004 (0.009)	0.005 (0.005)	-0.004 (0.008)
<i>N</i>	5,746		5,500		5,497	
AIC	225.600		-929.509		-1,518.802	
BIC	418.632		-737.746		-1,327.055	
Class probabilities	0.129	0.871	0.251	0.749	0.175	0.825
Class means	-0.027	0.298	-0.024	0.161	-0.054	0.197
Class σ	0.002	0.061	0.002	0.067	0.002	0.048

Notes: Estimates from model (7)-(8). Standard errors robust to heteroskedasticity and clustering at the household level. AIC: Akaike Information Criterion. BIC: Bayesian Information Criterion. Class probabilities: Estimates of π (class 1) and $1 - \pi$ (class 2).

Source: SLP v53, own calculations.

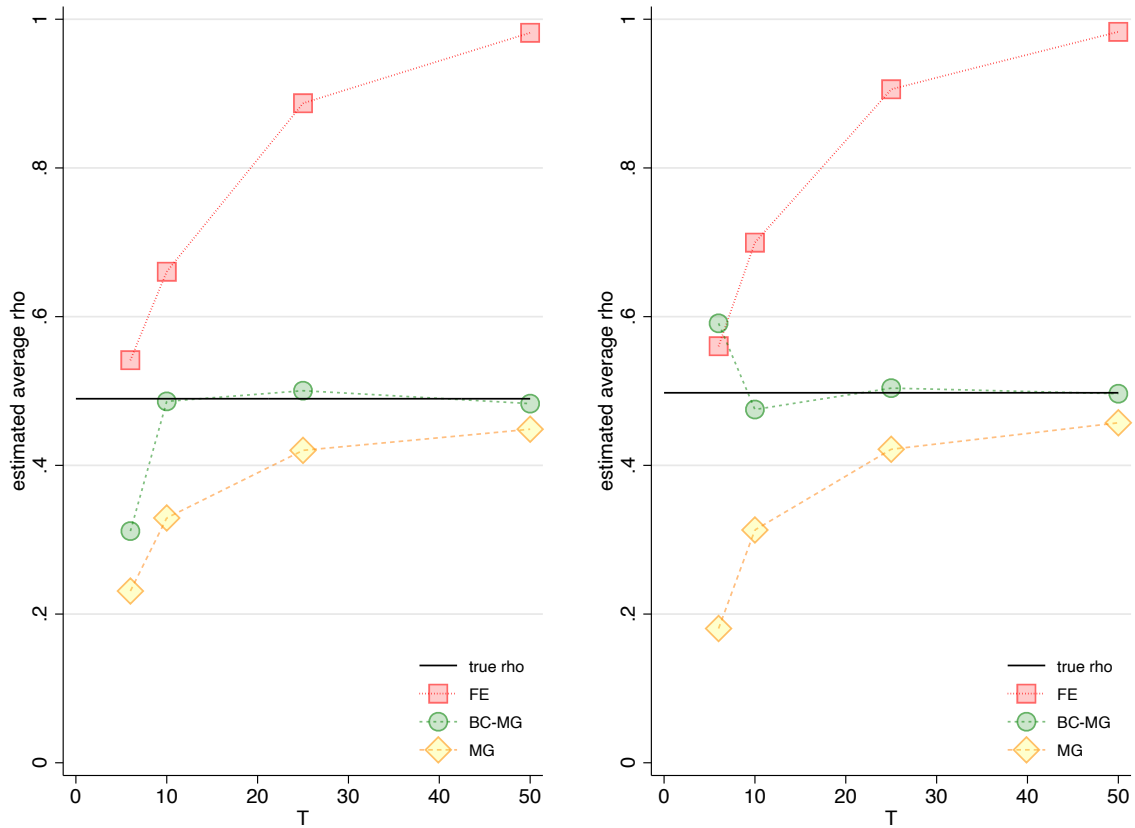


Figure 1: MONTE CARLO SIMULATIONS: ESTIMATED AVERAGE ρ_i OVER SAMPLES WITH DIFFERENT NUMBERS OF PERIODS (T)

Note: The figure plots average estimates of $E(\rho_i) = 0.5$ over 100 replications for the Fixed Effects estimator (“FE”), simple Mean Group estimator (“MG”) and bias-corrected Mean Group estimator (“BC-MG”) for the DGP described in the notes of Table 1 and where $\sigma_\rho = 0.25$. In the right-hand-side panel, 20 percent of observations have been deleted at random in each replication.

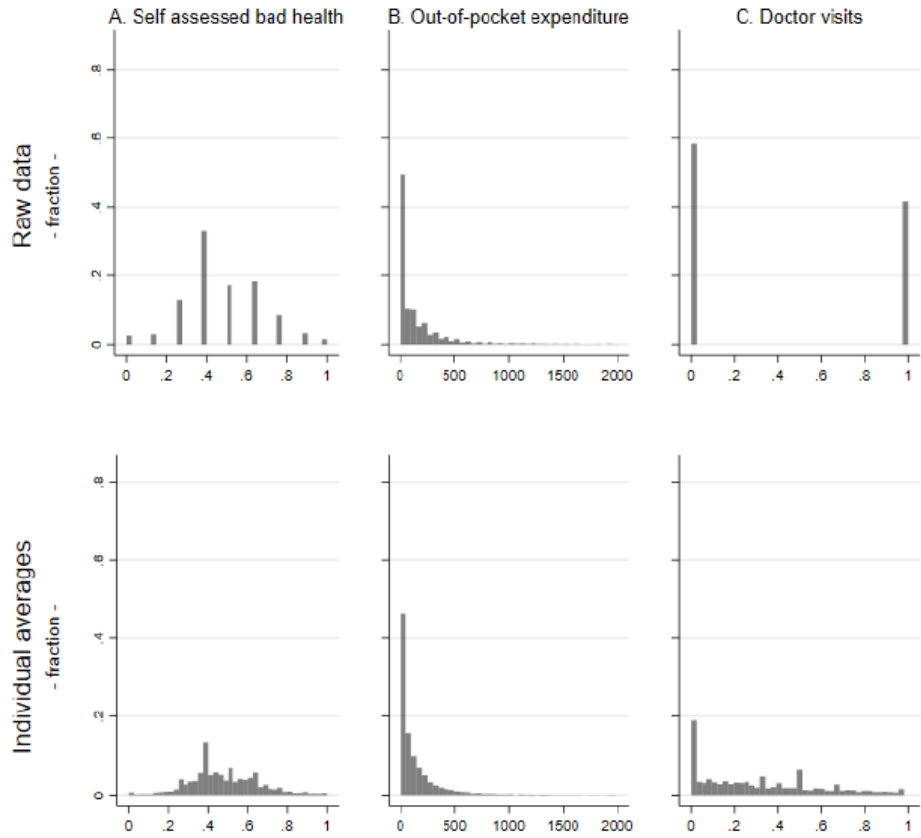


Figure 2: HISTOGRAMS OF OUTCOME VARIABLES AND THEIR PANEL MEANS

Note: Histograms of the outcome variables [top] and of their individual averages over time ($\bar{y}_i = \sum_{t=1}^{T_i} y_{it}/T_i$) [bottom].

Source: SLP v53, own calculations.

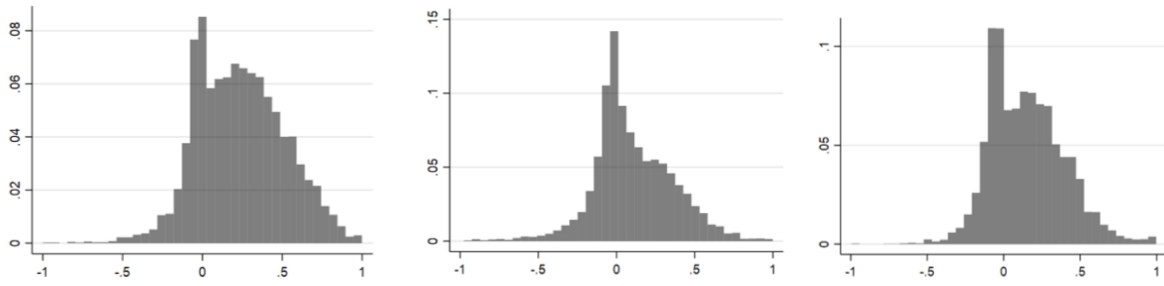


Figure 3: HISTOGRAMS OF ESTIMATED INDIVIDUAL-LEVEL STATE DEPENDENCE PARAMETERS $\hat{\rho}_i$

Note: Histograms of the estimated individual slopes ρ_i on the lagged dependent variable of model (2) corresponding to the estimates of Column (4) in Table 3 for each of the three outcomes self-assessed bad health (left figure), medical expenditures (middle) and doctor visits (right).

Source: SLP v53, own calculations.

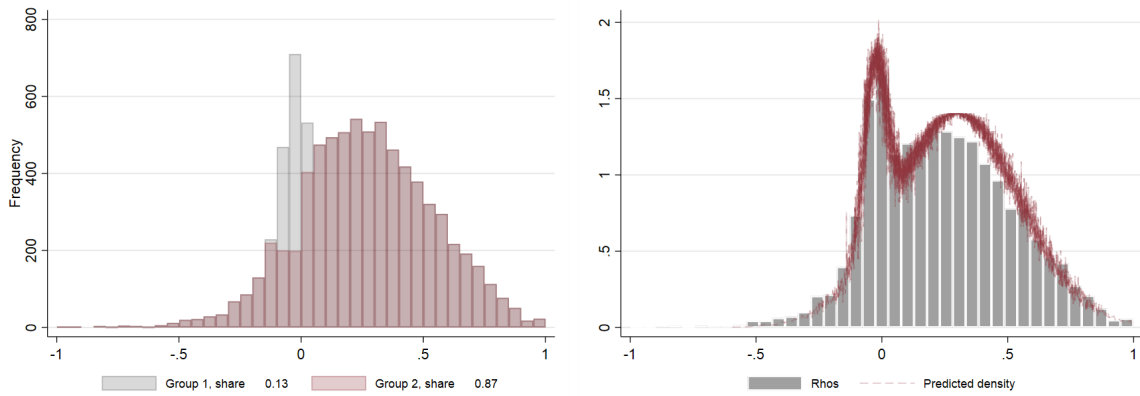


Figure 4: MODEL FIT — FINITE MIXTURE MODEL FOR PERSISTENCE IN SELF-ASSESSED BAD HEALTH, FULL SPECIFICATION

Note: Left-hand side: Histograms of predicted classes. Right-hand side: Point-wise predicted marginal density against histogram of $\hat{\rho}_i$. Graphs are based on estimates from Table 7.

Source: SLP v53, own calculations.

Appendix

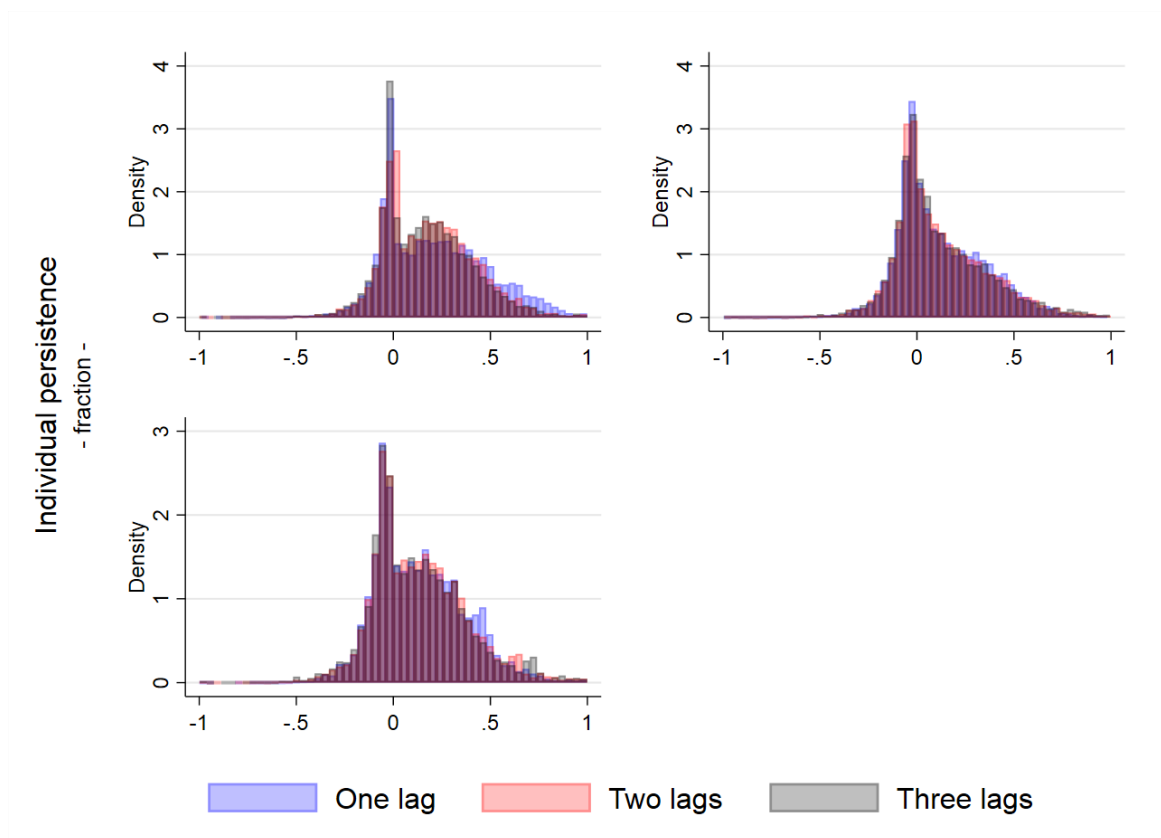


Figure A1: HISTOGRAMS OF ESTIMATED INDIVIDUAL-LEVEL STATE DEPENDENCE PARAMETERS $\hat{\rho}_i$, various lags

Note: Histograms of individual-specific coefficients on first lag of dependent variable from heterogenous AR (HAR) models. Blue for baseline HAR(1) model (2), identical to Figure 3; red for a HAR(2) model with first and second lag of the dependent variable (y_{it-1}, y_{it-2}); grey for a HAR(3) model with first, second and third lag of the dependent variable ($y_{it-1}, y_{it-2}, y_{it-3}$). Outcomes: Persistence in self-assessed bad health (top left), medical expenditures (top right), doctor visits (bottom).

Source: SLP v53, own calculations.

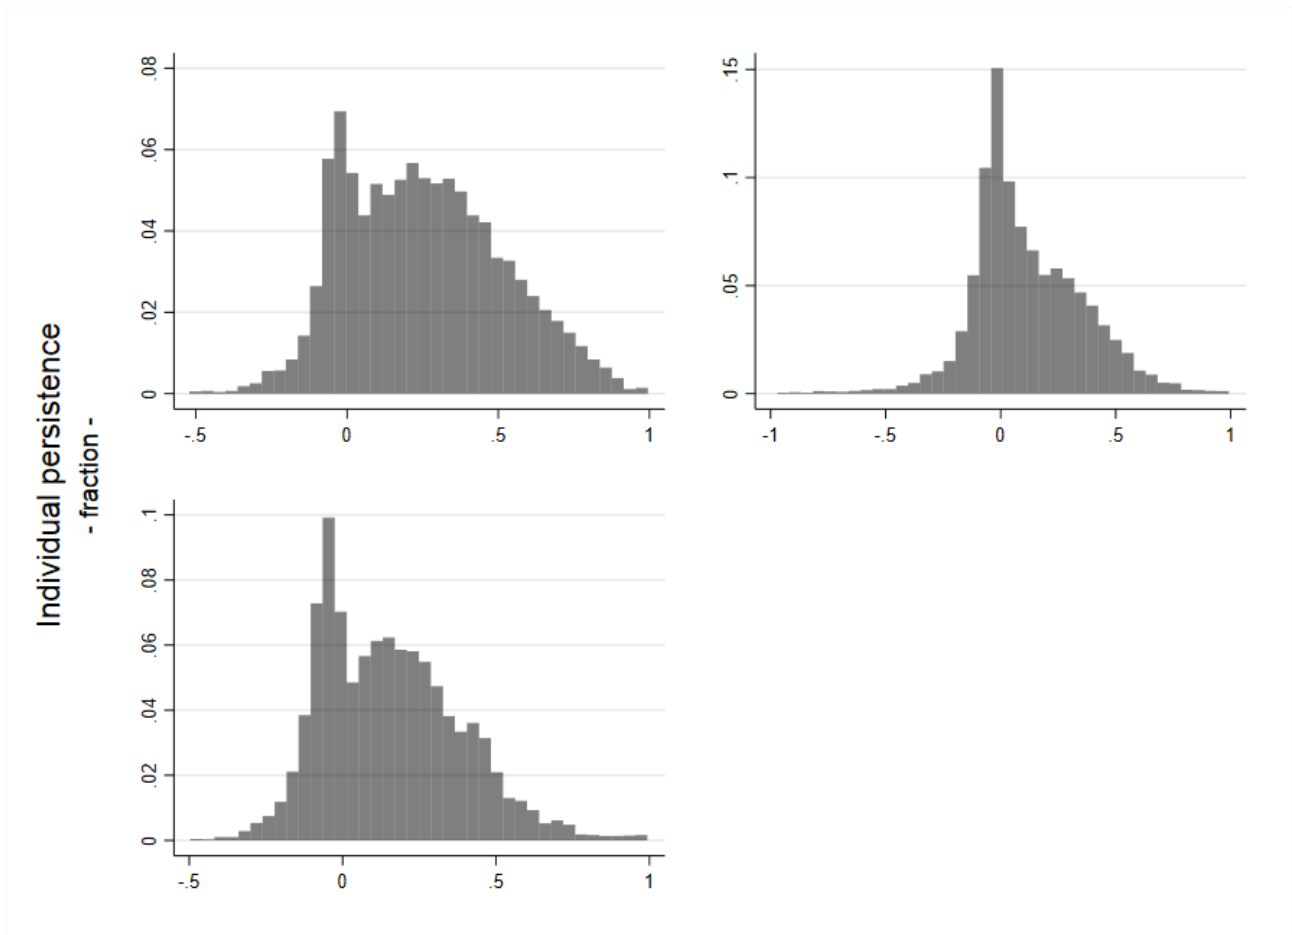


Figure A2: HISTOGRAMS OF ESTIMATED INDIVIDUAL-LEVEL STATE DEPENDENCE PARAMETERS $\hat{\rho}_i$, for $T > 30$

Note: Histograms analogous to Figure 3 but based only on $\hat{\rho}_i$ of individuals observed at least 30 times. Outcomes: Persistence in self-assessed bad health (top left), medical expenditures (top right), doctor visits (bottom).

Source: SLP v53, own calculations.

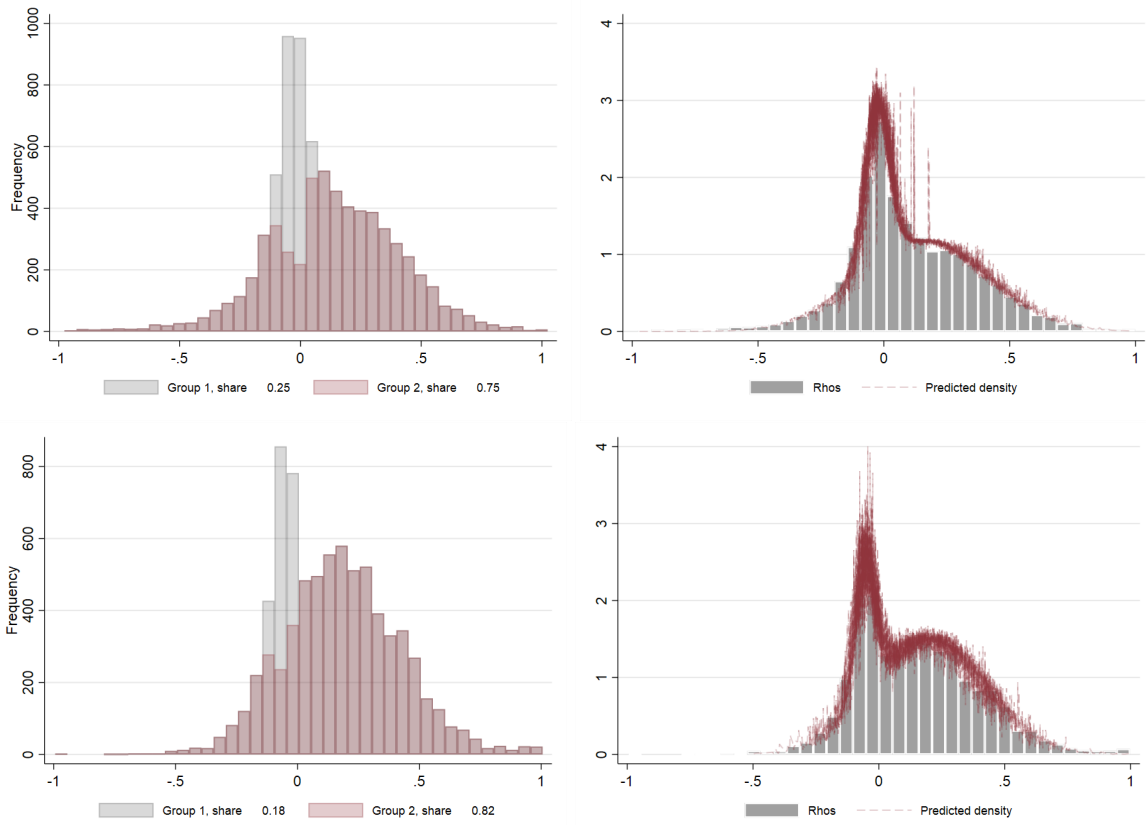


Figure A3: MODEL FIT — FINITE MIXTURE MODEL FOR PERSISTENCE IN MEDICAL EXPENDITURES AND DOCTOR VISITS, FULL SPECIFICATION

Note: Left-hand side: Histograms of predicted classes. Right-hand side: Point-wise predicted marginal density against histogram of $\hat{\rho}_i$. Top: Medical expenditures. Bottom: Doctor visits. All graphs based on estimates from Table 7.

Source: SLP v53, own calculations.