Bayesian learning in performance. Is there any?

Mike G. Tsionas*

July 10, 2022

Abstract

We propose and implement a Bayesian learning model for performance. The model implies a specific distribution for performance / technical inefficiency which we exploit in the context of stochastic frontier models. As the theoretical model is ambiguous with respect to what constitutes existing "experience", we propose and implement alternative specifications. The estimation and inference techniques are based on Bayesian analysis using Markov Chain Monte Carlo methods. We apply the new techniques to a data set of large U.S. banks. Our findings indicate that there is some learning in technical inefficiency although there is limited evidence, if at all, that jumps in experience are related to productivity growth. However, this effect is distinctly pronounced for the 2007-2010 period but much less significant afterwards.

Key Words: Productivity and Competitiveness; Production; Bayesian Learning; Bayesian methods; Organizational Structures.

*Montpellier Business & Lancaster University Management School, LA1 4YX, U.K., m.tsionas@montpellier-bs.com

1 Introduction

Despite the fact that performance measurement and estimation has received a great deal of attention in the literature using both Stochastic Frontier Analysis (SFA) and Data Envelopment Analysis (DEA) an important question has been left unanswered. Is it possible that decision makers over time may be able to learn the difficulties of the process, build experience and, eventually, reduce technical inefficiency? There are, of course, models where inefficiency depends on certain covariates or environmental variables and even models where inefficiency is persistent but these models are effectively "reduced forms" that do not allow for a systematic examination of learning in the process of generating efficiency of decision making units.

It is widely known that efficiency ultimately depends on managerial ability of managerial experience, a factor called "X-inefficiency" by Leibenstein. In the banking sector, for example, numerous studies have been conducted to measure or estimate technical or cost and profit efficiency using SFA and DEA.[1] But if there is systematic learning, of one form or another, and gradual building up of managerial experience, broadly defined, bank managers should be credited for being in the intermediate stages of learning. This is all the more important because cost inefficiency in banking has a dual role. First, it acts as a buffer stock which can be potentially reduced when non-performing loans build up as in the recent subprime crisis. For example, according to Ding and Sickles (2018) [27] "more efficient banks increase capital holdings and take on greater credit risk while reducing risk-weighted assets".

Second, cost inefficiency may be related to waste of resources that, however, provide for the "quite life" of a near monopolist (the term "quite life" is due to Hicks; see Koetter and Vins, 2008 [64]; Koetter et al., 2012 [63]). In the first case, we would expect nearly immediate reductions in inefficiency as a response to higher non-performing loans. In the second case, we would expect that in "normal times" no effort is made, on purpose, to reduce cost inefficiency and, as a result, no learning or experience and efficiency improvements will take place.

These traditional views may be challenged on the grounds that inefficiency, seen as waste of valuable resources, has an effect on several other performance indicators of a decision making unit (Delis et al., 2017 [24]). On these grounds, the management may try to learn the complex internal processes of the unit to reduce inefficiency. Of course, this process takes time and learning arises naturally in the form of learning about waste. [2] [3]

Learning or, more generally, entrepreneurial orientation and "ability" have been at the center of interest for analyzing performance (Ge and Peng, 2012; Stamboulis and Barlas, 2014; Ambad and Damit, 2016; Barba-Sanchez and Atienza-Sahuquillo, 2018). Especially, Moreno and Casillas (2008) and Rauch et al. (2009) focus on entrepreneurial orientation and performance. Moreno and Casillas (2008) emphasize the complexity of the relationships between entrepreneurial orientation, strategy, environment, resources and growth. This is not surprising, as learning involves activities along many dimensions such as ability to network, ability to assimilate experience and opportunities, ability to reflect on past experience, strategy and mistakes, ability to access resources and a component related to "entrepreneurial ability" (Deakins and Freel, 1998). See also Tortorella et al. (2020) and Azadegan et al. (2019).

When explicit data on learning are not available, the problem becomes more convoluted and learning has to be inferred indirectly from the data. For example, Xiao and Yang (2021) [98] propose a series of Bayesian learning and pricing models to uncover the unknown demand parameter through a data driven process (see also Wang and Xiao, 2017 [95]). Our problem in this paper is different, in the sense that *we allow for a flexible learning model to test formally whether learning is Bayesian, when only (indirect) production data are available*. The data is "indirect" in

[1] For bibliometric evidence on the use of SFA and DEA in banking see Lampe and Hilgers (2015) [65]. Papers that use DEA in banking include, *inter alia*, Andrew et al. (2018), Camanho and Dyson (2005) [71], Casu et al. (2016) [86], Fukuyama and Matusek (2017) [32], and Kevork et al. (2017) [62]. For papers that use SFA see Badunenko and Kumbhakar (2017) [85] and Tsionas et al. (2017 [88], 2020 [89]) as well as the references therein.

[2] Another factor is that learning may be costly see, for example, Ubï¿œe et al. (2017) [94] and Tsionas (2020) [89]. The issue is related to rational inefficiency, see Bogetoft and Hougaard (2013) [84] in the sense that costly improvements of efficiency are related to costly learning itself. This is formally examined in Tsionas and Mamatzakis (2017) [91] from a different perspective. Learning is, naturally, gradual, and may not be continual. Indeed, for the most part, managers may not worry too much about inefficiency unless it becomes a significant problem in some respects, in which case they may be *forced* to learn and improve their experience to reduce inefficiency. In the final analysis, although one can argue in this or the other way, the issue is, clearly, an empirical question.

Simple views of the learning process imply or assume that learning is a continuous process (Arrow, 1962 [82]; Young, 1993 [**?**], but see also Arrow, 2014 [83], 2016 [34]). Learning may depend on "intuition" or "mental models" (Johnson-Laird, 1983 [54], 10–11) or in formal models which acknowledge, of course, that they are imperfect representations of reality, a point made by Kahneman (2011) [58] and Mercier and Sperber (2017) [68] inter alia. Jaber et al. (2021) [53] suggest a learning curve that accounts "for the variable degree of cognitive interference that occurs while learning when moving from one repetition to the next". Learning curves are used widely in operations research and production operations management, see Jaber (2004, 2006 [49]), Jaber and Bonney (1999) [50], Jaber and Guifridda (2004) [51] Jaber and Khan (2010) [52], Fogliatto and Anzanello [30] (2011), Grosse et al. (2015) [45], Glock et al. (2019) [40], and Pusic et al. (2015 [78], 2017 [79]).

[3] There are many important applications of learning curves, in workforce planning (Cavagnini et al., 2020 [19]), machine scheduling (Anzanello et al., 2014 [23]), lot sizing (Jaber and Khan, 2010 [52]), Enterprise Resource Planning (ERP) implementations (Plaza and Rohlf, 2008 [77]), or order picking (Grosse and Glock, 2015 [45]).

the sense that learning is not directly (or even imperfectly) observed but has to be inferred from the available data. Related to learning are, of course, performance and efficiency of operations which are themselves unobservable and have to be estimated.

In this paper we adopt the view that the management applies Bayesian learning (which is quite natural for various forms and types of agents[4]) to build up experience and learn something about inefficiency. We are able to compare models where learning is gradual and we derive learning curves over time. Moreover, we allow and test for the possibility that learning is not continual but consists instead of jumps of unknown sizes at unknown time periods. The models are estimated using Bayesian techniques organized around Markov Chain Monte Carlo (MCMC) as well as the Particle Filtering (PF) algorithm (also known as Sequential Monte Carlo, SMC) for the more complex models with possible persistence and jumps in the learning process. The new techniques are applied to a data set of large U.S. banks and several interesting findings related to experience and learning are provided and discussed.

We provide extensive "reality checks" related to the benchmark learning model. Specifically, we compare the benchmark model with non-parametric versions of inefficiency dynamics to find that the benchmark performs roughly the same using as criteria both Bayes factors (also known as posterior odds ratios), dynamic behavior conditional on the environmental or contextual variables, steady-state or long-run inefficiencies, as well as a decomposition that splits inefficiency into learnable and non-learnable components. Moreover, we recognize explicitly that inefficiency estimates delivered by the data depend on information sets, more specifically the types of dynamic conditioning that we use. If we simply condition on the entire data set, the benchmark learning model delivers results that are inconsistent with general non-parametric formulations. However, when we condition on bank-specific information sets (possibly including summary statistics of information from other banks) the results are quite close. We proceed along these lines as managers may repeatedly process information using non-Bayesian heuristics (see Kahneman and Tversky, 1974 [59]; Camerer, 1998 [81]; Rabin, 1998 [80]). These studies contributed to advances in non-Bayesian learning (Golub and Jackson, 2010 [41]; Gilboa et al., 2008 [36], 2009 [37]) so, it is important to differentiate between different types of learning, if any, in the data. For example, there may be components in performance that can be learned and other components that cannot (relatively speaking) be learned. Moreover, when we process the data as a whole, it may not be possible to discriminate between the different types of learning that take place–where by "type of learning" we mean a particular conditioning on information that, however, uses Bayes' rule to update information. Other types of learning should, of course, be allowed and tested given the data.

We distinguish between passive and active learning, given the importance of these types of learning in the recent literature (e.g. Roy, Silvestre, and Singh, 2020) and we operationalize these concepts in models which are more general compared to our benchmark learning model. For similar operationalizations but from a different point of view, see Sabahi and Parast (2020). The basic implication is that, in the absence of direct data on learning, learning can be estimated more accurately when, after estimation, we condition on "reasonable" information sets rather than relying on mechanistic estimates of learning. The model passes successfully these different tests and, therefore, it can be considered as a learning model that reflects the broad picture of learning in large U.S. banks. Additionally, we consider decompositions of inefficiency into a component that is amenable to learning (and, therefore, potential reductions in waste) and a "non-learnable" component. The "non-learnable" component does not rely on simple residuals from a flexible model but rather, on post-processing which uses proper conditioning on the information sets that are likely to be faced by most bank managerial teams. To the best of our knowledge, the way this decomposition is implemented, as well as the decomposition itself, are novel, and show that, nearly, 20%-40% of observed inefficiency is non-learnable and can be attributed to the realities and complexities of the environment in which decisions are taken and learning takes place, while the remainder conforms to a relatively simple and, for the most part, standard model of learning.

2 Preliminaries

As noted in Minniti and Bygrave (2001): "[M]ost learning takes place by filtering signals obtained by experimenting with different competing hypotheses, where some actions are reinforced and others weakened as new evidence is obtained. Over time, we argue, individuals repeat only those actions that have generated better outcomes".

A decision maker must make a decision z related to an unknown quantity y. When $z = y$ then the decision maker attains maximum efficiency in the task. Following Jovanovic and Nyarko (1995) [57] we assume ability or efficiency (q) is given by

$$q = A\left[1 - (y - z)^2\right]. \tag{1}$$

[4]For example, in game theory in it standard to assume that agents follow Bayesian learning procedures (see, *inter alia*, Jordan, 1991 [56] and Nachbar, 1997 [72], 2005 [73]). For Bayesian analysis in the context of human rational learning, see Oaksford and Chater(2007) [74], and Kemp et al. (2007) [61].

After multiple trials (say T) the decision maker has

$$y_t = \theta + w_t,\ t = 1, \dots, T,\ w_t \overset{iid}{\sim} \mathcal{N}(0, \sigma_w^2) \tag{2}$$

where θ is a parameter, for example, an optimal way of performing tasks, optimal speed, optimal way of performing entrepreneurial-related activities etc. The entrepreneur has a prior on θ which is normal with a certain mean and variance σ_θ^2. It can be shown that Bayesian learning takes the following form

$$\mathbb{E}_T(q_T) = A(1 - x_T - \sigma_w^2), \tag{3}$$

where $\mathbb{E}_T$ denotes conditional expectation after T "trials", and

$$x_T = \frac{\sigma_w^2 \sigma_\theta^2}{\sigma_w^2 + T\sigma_\theta^2}, \tag{4}$$

is the posterior variance of θ, viz. $\mathbb{E}_T(\theta - \mathbb{E}_T\theta)^2$. The maximum value of (3) is $A(1-\sigma_w^2)$ so we can define a learning curve by

$$Q_T = \left(\frac{1 - x_T - \sigma_w^2}{1 - \sigma_w^2} \right), \tag{5}$$

where $Q_T \equiv \frac{\mathbb{E}_T(q_T)}{A(1-\sigma_w^2)^N}$. This curve will be S-shaped provided $\sigma_\theta^2 \geq \frac{2(1-\sigma_\theta^2)}{N+1}$.

3 Model

Suppose $x_{it} \in \mathbb{R}^K$ is a vector of log inputs, y_{it} is log output and production possibilities are given by

$$y_{it} = x_{it}'\beta + v_{it} - u_{it},\ i = 1, \dots, n,\ t = 1, \dots, T, \tag{6}$$

where $v_{it} \sim$ i.i.d$\mathcal{N}(0, \sigma_v^2)$ represents measurement error, and u_{it} is a non-negative random variable representing technical inefficiency. Following the Bayesian learning model of Jovanonic and Nyarko (1996) [57] and Oikawa (2016) [75], see also Tsionas (2017) [88], we denote $u_{it} = (\theta_{it} - z_{it})^2$ where z_{it} represents the manager's decision with respect to the inefficiency level, and θ_{it} is an unknown parameter. Therefore, the problem of the manager is to maximize expected efficiency, viz. $\mathbb{E}(e^{-u_{it}}) = \mathbb{E}\left[e^{-(\theta_{it} - z_{it})^2}\right]$ by choice of z_{it}.

If the manager knows θ_{it} or sets $\theta_{it} = z_{it}$ then production is fully efficient. So, "although a firm does not know how to optimize the utility of its own technology at first, it gradually learns better usage through production experience" (Oikawa, 2016 [75], p. 16). Suppose τ_{it} denotes firm experience, and

$$\theta_{it} \sim \text{i.i.d}\mathcal{N}(\bar{\theta}, \sigma_\theta^2),\ i = 1, \dots, n,\ t = 1, \dots, T, \tag{7}$$

is the prior of the critical parameter θ_{it}. Denote

$$\sigma(\tau) = \frac{\sigma_v^2}{\sigma_v^2/\sigma_\theta^2 + \tau} \tag{8}$$

In turn, Oikawa (2016) [75] showed that the distribution of u_{it} is *gamma*, viz. $\mathcal{G}\left(\frac{1}{2}, 2\sigma(\tau)\right)$, with density

$$f(u) = \frac{u^{-1/2} e^{-u/(2\sigma(\tau))}}{\sqrt{2\pi\sigma(\tau)}},\ u \geq 0, \tag{9}$$

see Greene (1980a [42], b [43], 1990 [44]). The mean of inefficiency is $\sigma(\tau)$ and its standard deviation is $\sqrt{2}\sigma(\tau)$. Defining $\lambda = \frac{\sigma_v^2}{\sigma_\theta^2}$ we have $\sigma(\tau) = \frac{\lambda\sigma_\theta^2}{\lambda+\tau}$.

Suppose $\theta = [\beta', \sigma_v, \lambda]'$, D denotes the data, and $p(\theta)$ is a prior on θ. The posterior distribution of the model is as follows:

$$\begin{aligned} p(\theta, \mathbf{u}|D) \propto \sigma_v^{-nT} \exp\left\{-\tfrac{1}{2\sigma_v^2} \textstyle\sum_{i=1}^n \sum_{t=1}^T \left(y_{it} + u_{it} - x_{it}'\beta\right)^2 - \tfrac{1}{2}\sum_{i=1}^n \sum_{t=1}^T \ln u_{it}\right\} \cdot \\ \exp\left\{-\tfrac{1}{2}\textstyle\sum_{i=1}^n \sum_{t=1}^T \frac{u_{it}}{\sigma(\tau_{it})} - \tfrac{1}{2}\sum_{i=1}^n \sum_{t=1}^T \ln\sigma(\tau_{it})\right\} \cdot p(\theta) \end{aligned} \tag{10}$$

where $\mathbf{u} = [u_{it};\ i = 1, \dots, n,\ t = 1, \dots, T]$, and $\sigma(\tau_{it}) = \frac{\lambda\sigma_\theta^2}{\lambda+\tau_{it}}$ Ideally, we would like to work with $p(\theta|D) =$

$_{\Re_{+}^{nT}} p(\theta,\mathbf{u}|D)\,\mathrm{d}\mathbf{u}$ but the integral is not available in closed form. Experience τ_{it} may be set to a time trend that is,

$$\tau_{it} = t. \tag{11}$$

Our prior has the form

$$p(\theta) \propto \sigma_v^{-1}\lambda^{-1}. \tag{12}$$

Writing out the prior in (12) explicitly, we obtain

$$\begin{aligned} p(\theta,\mathbf{u}|D) \propto \sigma_v^{-(nT+1)}\lambda^{-1}\exp\left\{-\tfrac{1}{2\sigma_v^2}\sum_{i=1}^{n}\sum_{t=1}^{T}(y_{it}+u_{it}-x_{it}'\beta)^2-\tfrac{1}{2}\sum_{i=1}^{n}\sum_{t=1}^{T}\ln u_{it}\right\}\cdot \\ \exp\left\{-\tfrac{1}{2}\sum_{i=1}^{n}\sum_{t=1}^{T}\tfrac{u_{it}}{\sigma(\tau_{it})}-\tfrac{1}{2}\sum_{i=1}^{n}\sum_{t=1}^{T}\ln\sigma(\tau_{it})\right\} \end{aligned} \tag{13}$$

Access to the posterior is provided by Markov Chain Monte Carlo (MCMC) techniques described in part A.1 of the Technical Appendix. We call this specification **Model I**.

4 Alternative models for experience

In the previous section, we assumed τ_{it} is a simple time trend, which involves the assumption that experience simply accumulates with the passage of time. Although this may well be correct we would like, nevertheless, to test such hypotheses, in the light of the data, using alternative specifications for experience. Our first alternative (which we call **Model II**) is a latent variable autoregressive scheme for experience:

$$\tau_{it} = \alpha_i + \rho\tau_{i,t-1} + \xi_{it},\ \xi_{it} \sim \text{i.i.d}\,\mathcal{N}(0,\sigma_\xi^2), \tag{14}$$

where α_i denotes individual effects, ρ is an autoregressive parameter, and ξ_{it} is an error term. For the firm effects we assume $\alpha_i \sim \text{i.i.d}\,\mathcal{N}(\bar{\alpha},\sigma_\alpha^2)$. A more reasonable model, which we call **Model III**, *contains a jump component to allow for the possibility that experience possibly increases only at discrete time intervals*:

$$\tau_{it} = \alpha_i + \rho\tau_{i,t-1} + J_{it} + \xi_{it},\ \xi_{it} \sim \text{i.i.d}\,\mathcal{N}(0,\sigma_\xi^2), \tag{15}$$

where J_{it} is zero with probability ϖ_{it}, otherwise it assumes an unknown value γ_{it} with probability $1-\varpi_{it}$. The interpretation that experience possibly increases only at discrete time intervals involves the restriction $\rho = 0$ although we can have intermediate cases where $\rho \neq 0$ and there is a jump component, simultaneously.
Alternatively, we can write (15) as follows.

$$\tau_{it} = \alpha_i + \rho\tau_{i,t-1} + \gamma_{it}\mathscr{J}_{it} + \xi_{it},\ \xi_{it} \sim \text{i.i.d}\,\mathcal{N}(0,\sigma_\xi^2), \tag{16}$$

where $\mathscr{J}_{it}$ is zero with probability ϖ_{it}, and equal to one with probability $1-\varpi_{it}$. We assume that jump sizes come from a distribution $\mathcal{N}(\gamma_o,\sigma_\gamma^2)$ where γ_o and σ_γ are unknown parameters. For the prior parameters, we assume $\gamma_o \sim \mathcal{N}(0,\bar{s}^2)$ and $\frac{\bar{q}}{\sigma_{\gamma^2}} \sim \chi^2_{\bar{n}}$, i.e. an inverted *gamma* distribution. This prior is minimally informative if we set $\bar{q} = 0.01$, $\bar{n} = 1$ and $\bar{s} = 2$. Notice that the jump size can be positive or negative.
Allowing for jumps in experience is important in view of the evidence that unchanged environments do not provide much opportunity for learning whereas diverse experiences contribute to solid increases in entrepreneurial orientation (see, for example, Deakins and Freel, 1998, p. 151; Azadegan et al., 2019).
For the probabilities we assume the following process:

$$\ln\tfrac{\varpi_{it}}{1-\varpi_{it}} \sim \text{i.i.d}\,\mathcal{N}(r_o,\sigma_\varpi^2), \tag{17}$$

where $\ln\frac{\varpi_{it}}{1-\varpi_{it}}$ is the log odds ratio, r_o is its prior mean, and σ_ϖ its prior standard deviation, both of which are unknown parameters. For various parameterizations of the prior in (17) the resulting priors in terms of ϖ are presented in Figure 1.

As σ_ϖ increases the prior becomes more diffuse and r_o mainly controls the probability density around the endpoints (one and zero). From this evidence, it seems reasonable to select $r_o = 0$ and $\sigma_\varpi = 10$ to produce a relatively diffuse prior for ϖ. Nevertheless, we prefer to impose this prior information in a stochastic way to allow that r_o and σ_ϖ

Figure 1: Implied priors of probability parameter, ϖ

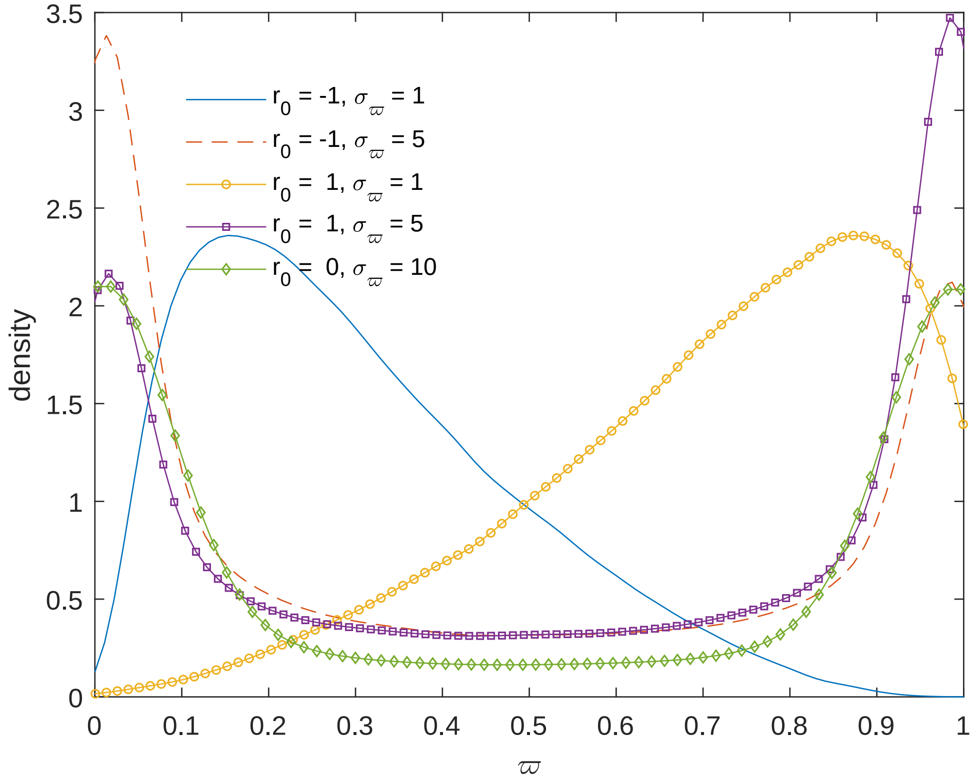

Table 1: Models

Model	name	equation
I	Time trend model	(11)
II	Latent AR model	(14)
III	Latent AR model with jumps	(15)

are unknown parameters, via the following form:

$$r_o \sim \mathcal{N}(0,1),\ p(\sigma_{\varpi}) \propto \sigma_{\varpi}^{\bar{N}-1} e^{-\bar{N}\bar{Q}/(2\sigma_{\varpi}^2)}, \tag{18}$$

viz. r_o follows a standard normal distribution, and σ_{ϖ} follows a variant of the inverted *gamma* distribution (Zellner, 1971 [?], p. 371). If we set the shape parameter $\bar{N} = 2$ then the mode of the prior for σ_{ϖ} is 10 if we set the scale parameter $\bar{Q} = 10\sqrt{\frac{3}{2}} \simeq 12.2$. Finally, as the prior could be important we undertake extensive sensitivity analysis in part A.3 of the Technical Appendix. For Models II and III, the Gibbs sampler could be numerically inefficient (for a precise notion see also part A.3 of the Technical Appendix) so, we apply a Particle Filtering (PF) approach also known as Sequential Monte Carlo (for details see part A.2 of the Technical Appendix). In all cases, we use 250,000 iterations the first 50,000 of which are discarded in the interest of mitigating possible start up effects.
To summarize the models, see Table 1.

5 Data and empirical results

We have an unbalanced panel with 3,897 bank-year observations for 285 large U.S. commercial banks operating in 2001-2019, whose total assets were in excess of one billion dollars (in 2005 U.S. dollars) in the first three years of observation. The data come from Call Reports available from the Federal Reserve Bank of Chicago and they have been used in Malikov et al. (2016) [67] but have been updated to include data for the period 2011-2019. For detailed description of the data construction, see Section 5 of the paper.
The list of included variables is as follows: y_1 Consumer Loans, y_2 Real Estate Loans, y_3 Commercial & Industrial Loans, y_4 Securities, y_5 Off-Balance Sheet Activities Income, b Total Reported Nonperforming Loans, x_1 Labor, number of full-time employees, x_2 Physical Capital (Fixed Assets), x_3 Purchased Funds, x_4 Interest-Bearing Transaction Accounts, x_5 Non-Transaction Accounts, e (Quasi-Fixed) Equity Capital.[5]
Suppose $x_{it} = [x_{it,1}, \ldots, x_{it,K}]'$ denotes the logs of inputs and $y_{it} = [y_{it,1}, \ldots, y_{it,M}]'$ is the vector containing the logs of outputs. To model the technology we use an output distance function (ODF). To this end, we define $\tilde{y}_{it,m} = y_{it,m} - y_{it,1}$ $(m = 2, \ldots, M)$. The output distance frontier is as follows.

$$y_{it,1} = \beta_o + \beta_1' z_{it} + \tfrac{1}{2} z_{it}' B z_{it} + v_{it} - u_{it}, \tag{19}$$

where $z_{it} = [\tilde{y}_{it}', x_{it}']'$, $\tilde{y}_{it} = [y_{it,2}, \ldots, y_{it,M}]'$, β_o is a scalar, β_1 is a vector and B is a symmetric matrix all containing unknown parameters which we denote $\beta = [\beta_o, \beta_1', vech(B)]'$ where $vech$ is the operator that stacks different elements of a matrix into a vector.[6]
Model comparison is facilitated by the use of marginal likelihoods and Bayes factors (Kass and Raftery, 1994 [60]). For any model with parameters θ_j and data D, the marginal likelihood is the integrating constant of the posterior, i.e.

$$\mathscr{M}(D) = \mathscr{L}(\theta; D) p(\theta) \mathrm{d}\theta, \tag{20}$$

where $\mathscr{L}(\theta; D), p(\theta)$ denote, respectively, the likelihood and the prior. For any two models indexed by "j" the Bayes factor in favor of model "j" and against model "1" is the ratio of marginal likelihoods:

$$BF_{j:1} = \tfrac{\mathscr{M}_j(D)}{\mathscr{M}_1(D)}, \tag{21}$$

[5] Unless otherwise noted all variables are in thousands of U.S. dollars.
[6] The output distance function is non-increasing in outputs and non-decreasing in inputs. We impose these properties at the means of the data and ten other randomly selected points. In turn, we check whether estimated (posterior mean) input and output elasticities satisfy these properties. Equity over total assets and non-performing loans (b) are included in the 19. Non-performing-loans have derivative properties similar to inputs while no restrictions are imposed on equity over total assets.

Table 2: Bayes factors against Model I

	Model II	Model III
2005	26.71	4,755.2
2006	17.19	9.821.3
2007	22.71	12,781.4
2008	48.83	19,655.0
2009	57.20	22,455.7
2010	44.50	24,718.8

Figure 2: Aspects of the model

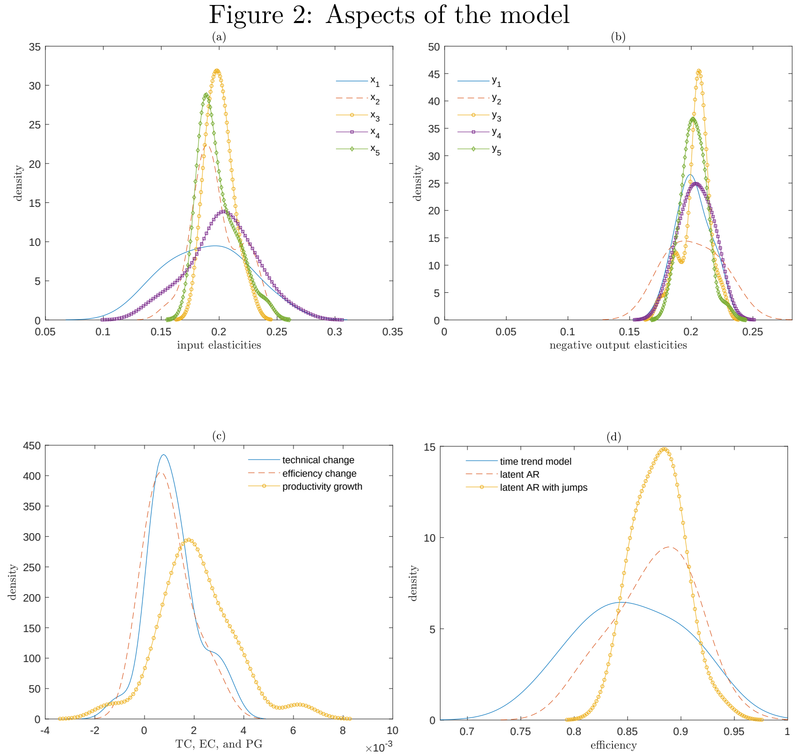

$(j = 2, 3, \ldots, J)$ which is also known as posterior odds ratio when the prior odds ratio is 1:1.[7] Here, J denotes the number of models under consideration.

For Model I (time trend) the marginal likelihood can be computed using the Laplace approximation (DiCiccio et al., 1997 [26]). For Models II and III it is standard output of the Particle Filtering (PF) approach. To compare the different models, we introduce temporal observations in batches. First, we use all observations for 2000-2005, then we introduce all observations for 2006, etc. Bayes factors reported in Table 2, show that the odds in favor of Model III are overwhelming suggesting that, in the light of the data, this model receives the greatest support. The evidence in favor of Model II is also significant (although not as overwhelming as in the case of Model III). Bayes factors in favor of Model III and against Model II can be computed as $BF_{3:2} = \frac{\mathscr{M}_3(D)}{\mathscr{M}_3(D)} = \frac{\mathscr{M}_3(D)/\mathscr{M}_1(D)}{\mathscr{M}_2(D)/\mathscr{M}_1(D)}$. For example, in 2009, the Bayes factor in favor of Model III and against Model II is, approximately, 22,455.7/57.20=392.6, which is substantial. Given the superiority of Model III all results (unless otherwise noted) refer to Model III.

Figure 3: Marginal posterior densities of σ_θ, σ_v, σ_τ, and ρ

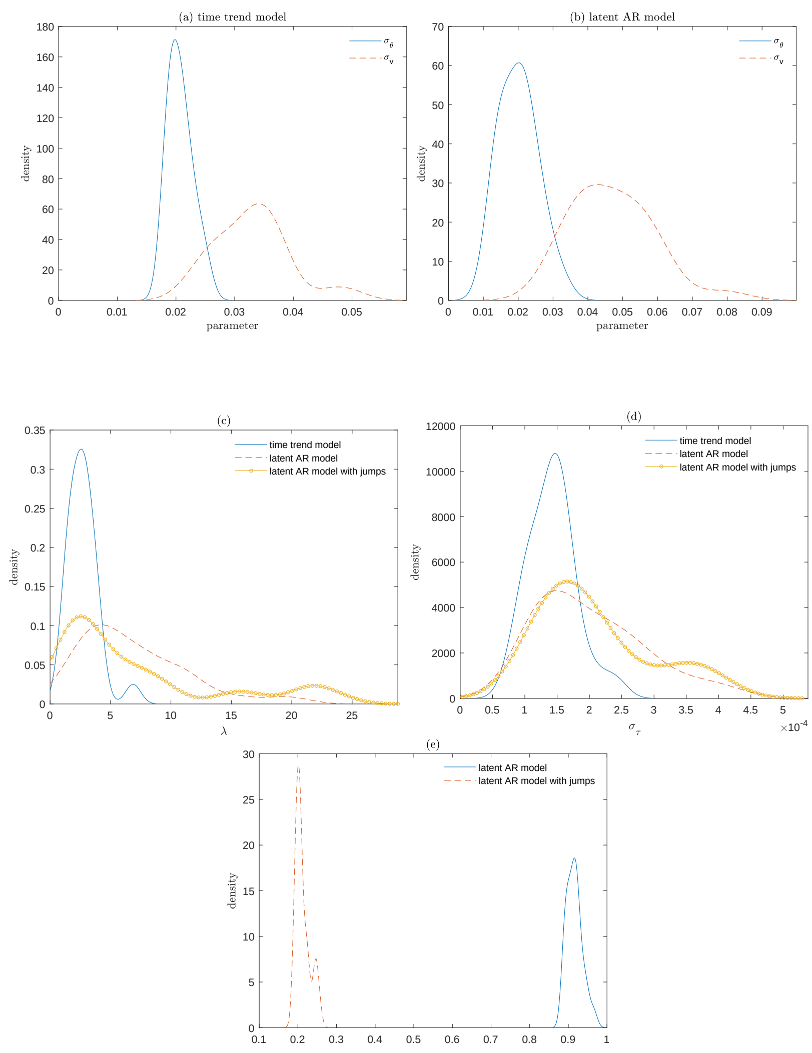

Figure 4: Posterior expectation of τ_{it} and 95% HPD intervals

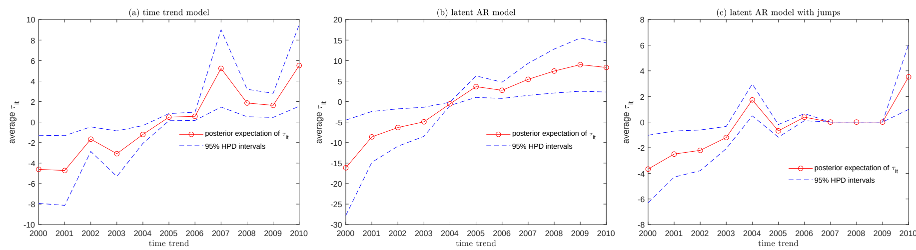

Notes: "HPD" stands for highest posterior density.

Figure 5: Normalized learning curves

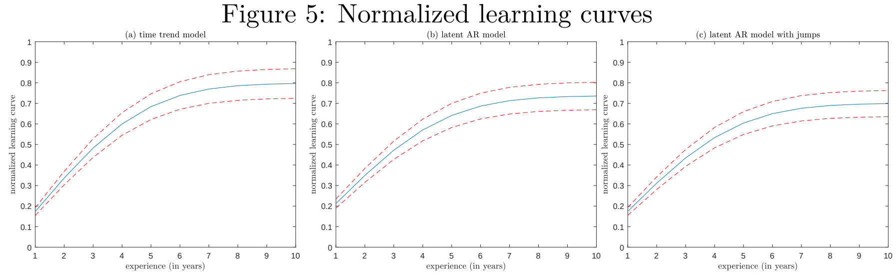

Notes: "HPD" stands for highest posterior density.

Finally, we report normalized learning curves in Figure 5. Normalized learning curves have been proposed by Jovanovic and Nyarko (1996). *The evidence suggests that learning is quite slow* as moving from an efficiency level of 0.6 to 0.7 would take, on the average, close to five years.

In Figure 6 we report the marginal posterior densities of ϖ_{it} in panel (a) and the jump parameter γ_{it} in panel (b).

The jump probability averages 0.188 (with posterior standard deviation 0.096) showing that the probability of learning taking place is rather low. The posterior mean of the jump size, γ_{it}, is 0.285 (with posterior standard deviation 0.036).
As the jump probability and size are time-varying and bank-specific, it is interesting to examine their relationship (using the posterior means of ϖ_{it} and γ_{it}). For example, a positive (negative) relationship would indicate that higher jump probabilities are associated with higher (lower) jump sizes. Aspects of the bivariate posterior distribution are reported in panels (c) and (d) from which there exists a positive relationship (the correlation coefficient is 0.64 but the relationship is unlikely to be linear as the bivariate posterior distribution is not unimodal).
An important assumption in the model is that the jump probability and size are not persistent, that is ϖ_{it} and γ_{it} do not depend on their lagged values ($\varpi_{i,t-1}$ and $\gamma_{i,t-1}$). To test this assumption we provide, in Figure 7, aspects of their bivariate posterior distributions.

From this evidence we have no reason to doubt that there is no persistent element in ϖ_{it} and γ_{it}. Moreover, from the evidence in the bivariate posterior distributions in Figure 8, we fail to uncover a relationship between the jump probability and jump size on the one hand, and posterior mean productivity growth on the other.

[7] It is relevant that marginal likelihoods and Bayes factors are not unimportant in discussions about causal structure. Fong and Holmes (2020) [31], for example, showed that the marginal likelihood is formally equivalent to exhaustive leave-p-out cross-validation averaged over all values of p and all held-out test sets when using the log posterior predictive probability as the scoring rule." (Fong and Holmes, 2020 [31], p. 489).

Figure 6: Marginal posterior densities of ϖ_{it} and γ_{it}

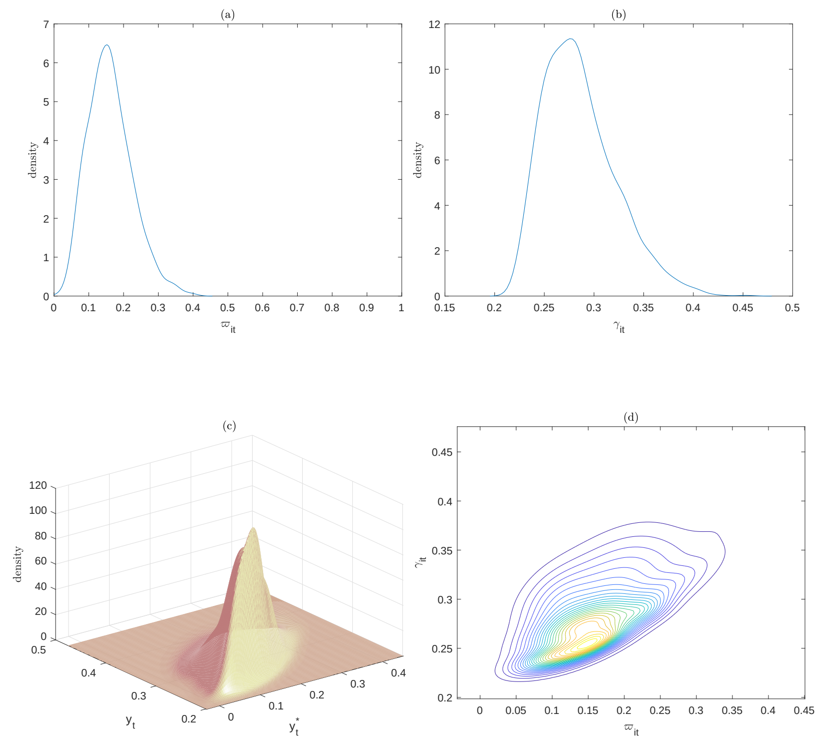

Figure 7: Aspects of the bivariate posterior distributions

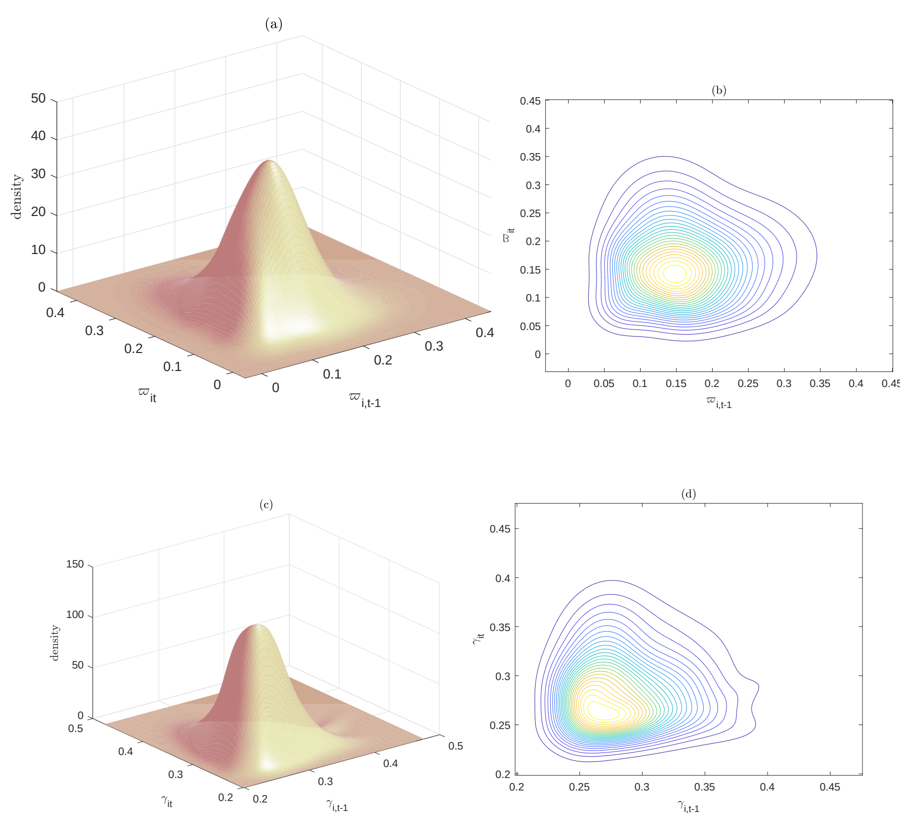

Figure 8: Aspects of the bivariate posterior distributions for productivity growth

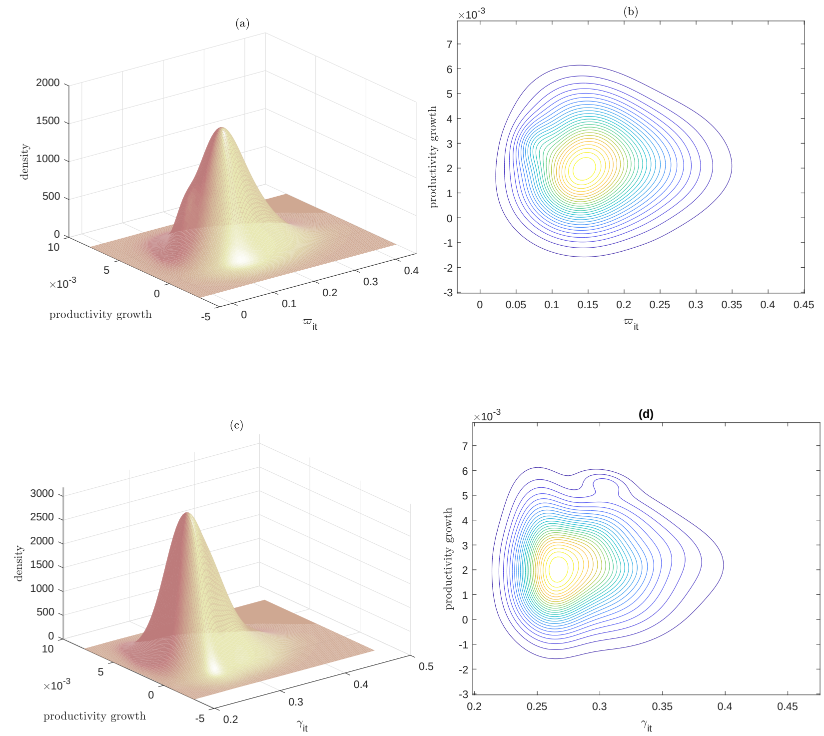

Therefore, jumps in experience do not seem to be related necessarily to low productivity growth per se. Some further evidence is provided in Figure 9, where we report the marginal posterior distributions of jump sizes by year (across all banks).

From the evidence in Figure 9, it turns out that during 2000-2006 there is very little evidence of jumps in experience (panel (a)) although there is more evidence in favor of improvements in experience during the subprime crisis (2007-2010, panel (b)). Results for the period 2011-2019 are reported in panel (c). From this evidence, it is clear, that jump sizes are fewer and estimated to be between 0.015 and 0.02 with a median close to 0.018.
To investigate further these results, we examine the relationship between jump sizes and jump probabilities (as in Figure 8) for the period 2007-2010. Our results are reported in Figure 10.

The evidence suggests a relatively strong positive relationship between jump probability and productivity growth

Figure 9: Marginal posteriors of jump size by period

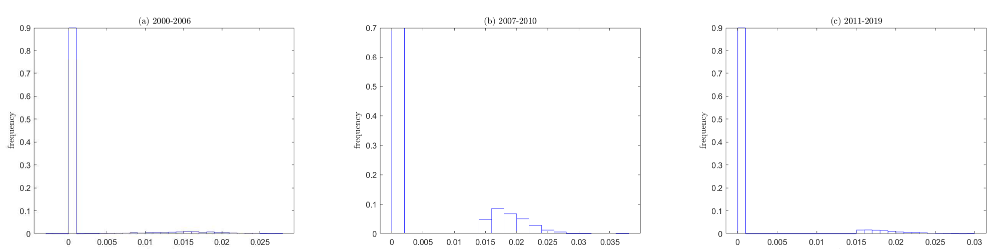

Figure 10: Aspects of the bivariate posterior distributions (2007-2010)

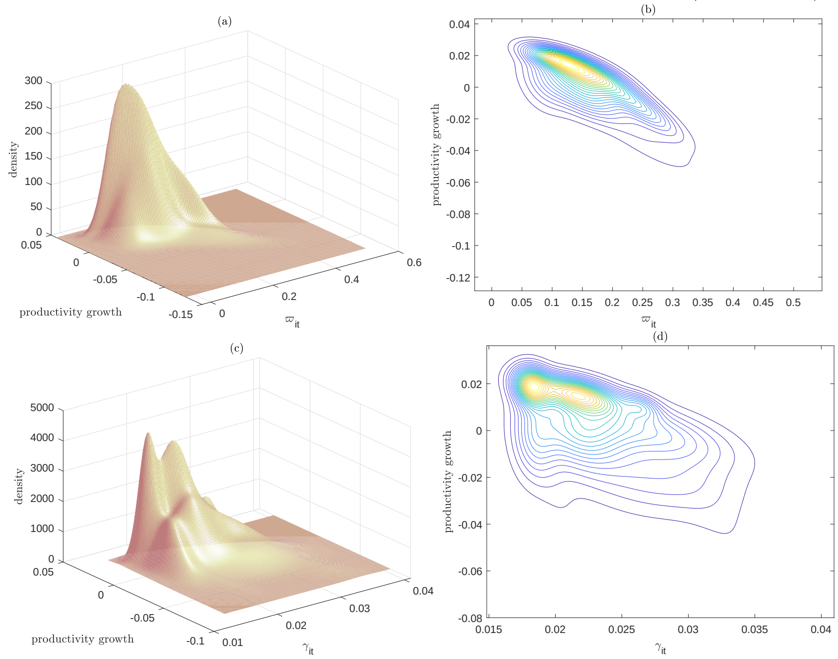

(panels (a) and (b)). The relationship between jump size and productivity growth (panels (c) and (d)) is also positive although the evidence suggests that the bivariate posterior is bimodal and there may be different relationships around the different modes or clusters in the data. The correlation coefficients between productivity growth and ϖ_{it}, γ_{it} are close to –0.77 and –0.56, respectively. As the relationship is negative and much more pronounced, in effect, this shows that lower productivity growth increases the probability of a jump and the jump size in experience itself.
Another important issue is whether non-performing loans (NPL) have a role in learning and experience. In principle, mounting NPLs may force banks to adopt different risk management procedures (Boussemart et al., 2019 [39]) but also to reduce inefficiencies to save on valuable resources. We examine this matter in Figure 11, for three subperiods, 2000-2007, 2007-2010, and 2011-2019, that is, roughly, before and after the subprime crisis. For visual clarity we normalize logs of NPLs between zero and one. Marginal posteriors are reported in Figure 11. From the evidence it turns out that before the subprime crisis, despite an overall positive relationship between NPLs and experience there are, in fact, three clusters or business models in the data, and in all of them the relationship between the two variables is rather weak (see panels (a) and (b)). After the financial crisis, there are still three clusters but the relationship becomes positive suggesting reinforcement of learning or experience in view of the increase in NPLs.

6 Reality checks

Models and Bayesian learning rely on strong assumptions as it is well known. For the underlying issues, see Jones and Love (2011). As they mention: "In the context of rational Bayesian modeling, existence proofs hide the fact that there are generally many Bayesian models of any task, corresponding to different assumptions about the learner's goals and model of the environment. Comparison among alternative models would potentially reveal a great deal about what people's goals and mental models actually are. Such an approach would also facilitate comparison to models within other frameworks, by separating the critical assumptions of any Bayesian model (e.g., those that specify the learner's generative model) from the contribution of Bayes' Rule itself. This separation should ease recognition of the logical relationships between assumptions of Bayesian models and of models cast within other frameworks, so that theoretical development is not duplicated and so that the core differences between competing theories can be identified and tested".
For example, Hautsch and Hess (2007) [47] use measures of quality of released information in financial markets and trading decisions. On the basis of precision proxy variables, they document significant evidence in favor of the

Figure 11: Relationship between learning and NPLs

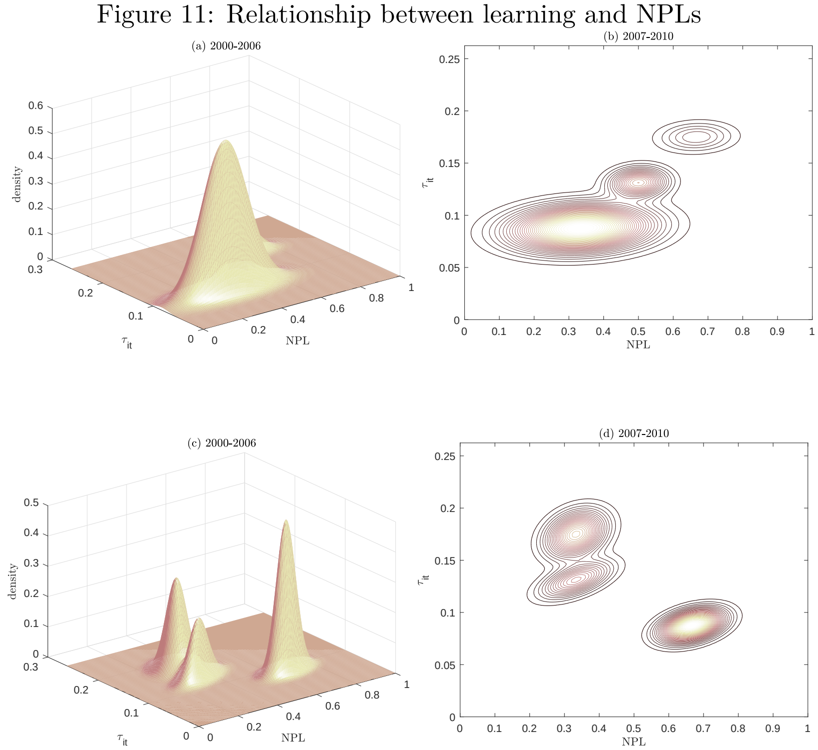

hypothesis of Bayesian learning (viz. that the quality of information acts as a catalyst, i.e., prices respond stronger to more precise news). Moreover, a series of experiments suggest that people may repeatedly process information using non-Bayesian heuristics (see Kahneman and Tversky, 1974 [59]; Camerer, 1998 [81]; Rabin, 1998 [80]). These experiments contributed to advances in non-Bayesian learning (Golub and Jackson, 2010 [41]; Gilboa et al., 2008 [36], 2009 [37]).

To compare the particular Bayesian learning model in this study with a general alternative model, the latter must allow for arbitrary types of learning. Such a general model is based on insights from Paul and Shankar (2018) [76] and Tsionas and Mamatzakis (2019) [92]. Since efficiency is $r_{it} = e^{-u_{it}}$ it follows that efficiency is bounded between zero and one so, any distribution function can be used to parametrize explicitly efficiency as a function of contextual or environmental variables. For example, if such variables are in vector $w_{it} \in \mathbb{R}^{d_w}$ one can assume

$$r_{it} = F\left(w_{it}; \delta\right), \tag{22}$$

where $F(\cdot)$ is any distribution function (for example, the standard normal or logistic), and $\delta \in \mathbb{R}^{d_\delta}$ is a vector of parameters. From (22) it follows that:

$$u_{it} = -\log F\left(w_{it}; \delta\right). \tag{23}$$

Fundamental to every type of learning is, of course, the presence of dynamics in technical inefficiency so, we have to modify the model as

$$u_{it} = -\log F\left(u_{i,t-1}, w_{it}; \delta\right), \tag{24}$$

keeping the same notation for parameters δ in the interest of simplicity. To operationalize (24), we use a flexible model:

$$u_{it} = -\log \sum_{g=1}^{G} \delta_{(g),1} \Phi\left(\delta_{(g),2} u_{i,t-1} + w_{it}' \delta_{(g),3}\right) + \xi_{it}^{u}, \tag{25}$$

where $\Phi(a) = \frac{1}{1+e^{-a}}$ $(a \in \mathbb{R})$ is the logistic distribution function, and ξ_{it}^{u} is an error term which is absent from Paul and Shankar (2018) [76] and Tsionas and Mamatzakis (2019) [92]. This error term may be critical in capturing the effects of omitted variables. The representation in (25) is, essentially, an artificial neural network (ANN) whose

Table 3: Bayes factors and posterior model probabilities

Bayes factors				PMP			
$G = 1$	1.000	$G = 6$	12.43	$G = 1$	0.0012	$G = 6$	0.015
$G = 2$	144.21	$G = 7$	5.51	$G = 2$	0.178	$G = 7$	0.007
$G = 3$	255.30	$G = 8$	1.30	$G = 3$	0.316	$G = 8$	0.002
$G = 4$	212.423	$G = 9$	0.043	$G = 4$	0.263	$G = 9$	0.000
$G = 5$	175.83	$G = 10$	0.001	$G = 5$	0.218	$G = 10$	0.000

good approximation properties are well-known (White, 1989 [96], 1990; Hornik et al., 1989 [48]). In (25), $\delta_{(g),1} > 0$ and $\delta_{(g),2}$ are scalar parameters, and $\delta_{(g),3} \in \mathbb{R}^{d_w}$ (for all $g = 1, \ldots, G$).
We proceed under the assumption that

$$\xi^u_{it} \overset{iid}{\sim} \mathcal{N}(0, \sigma^2_{\xi^u}). \tag{26}$$

In turn, our purpose is to estimate jointly (19) and (25). Implementation of MCMC using SMC is based on 150,000 iterations omitting the first 50,000 in the interest of mitigating possible start up effects. Denoting

$$\delta = [\delta_{(g),1}, \delta_{(g),2}, \delta'_{(g),3}]',$$

our prior for (25) is flat:

$$p(\delta, \sigma_{\xi^u}) \propto \sigma^{-1}_{\xi^u}. \tag{27}$$

As marginal likelihood is standard output of SMC, we select the order GN by fitting models with different values of G (ranging from 1 to 10). In turn, we can compute Bayes factors as in (21) and so-called posterior model probabilities (PMP) defined as probabilities that a particular model is "correct" in the light of the data:

$$\mathscr{P}_j(D) = \frac{\mathscr{M}_j(D)}{\sum_{j'=1}^{J} \mathscr{M}_{j'}(D)},\ j = 1, \ldots, J. \tag{28}$$

where marginal likelihoods, $\mathscr{M}_j(D)$, are defined in (20). For any function of the data and the parameters, say $q(\theta, D)$, the BMA posterior expectation of this function of interest is

$$\mathbb{E}_{\theta|D}\left[q(\theta, D)\right] = \sum_{j=1}^{J} \mathscr{P}_j(D) q(\theta_j, D), \tag{29}$$

where $q(\theta_j, D)$ is the function of interest for a particular model j, and $\mathbb{E}_{\theta|D}\left[\cdot\right]$ denotes expectation with respect to the parameters given the data. The results are reported in Table 3.

Normalizing the Bayes factor corresponding to $G = 1$ to unity, models with G =2, 3, 4, and 5 account, collectively, for, nearly, 97.5% of posterior model probability. In this instance, it does not seem wise to select a particular model (say, the model with the highest PMP, viz. $G = 3$) but rather to perform Bayesian Model Averaging (BMA) for different orders of G in (25).
We are especially interested in the function of interest that represents the learning rate

$$q_{\text{passive}}(\theta, D) = \frac{\partial \mathbb{E}_{\theta|D}(u_{it})}{\partial u_{i,t-1}}, \tag{30}$$

resulting from (25) and (29). We call this the **passive learning approach** as it is based on the entire data set, and represents an ex post (or a posteriori) overall estimate of learning. The passive learning estimate, relies on data from all banks and all time periods so, it is not unlike fitting learning curves to learning data.
The active learning approach is based on the idea that learning effects may be captured more accurately if we allow for sequential processing of data which is what the analyst and the decision maker share in common when the decision makers learns how to perform better. Suppose $\mathcal{I}_{i,t-1}$ denotes the information set consisting of data up to and including period $t - 1$ (i.e. the beginning of period t, in which decisions about inputs, outputs and results for performance are not yet known both to the researcher and the decision maker). The active learning estimate that is equivalent to (30) becomes

$$q_{\text{active},1}(\theta, D) = \frac{\partial \mathbb{E}_{\theta|\mathcal{I}_{i,t-1}}(u_{it})}{\partial u_{i,t-1}}. \tag{31}$$

The conditional expectation $\mathbb{E}_{\theta|\mathcal{I}_{i,t-1}}(\cdot)$ relies on estimating (19) and (25) in a different way, as we do not condition on the entire data set. Specifically, we are interested in posterior distributions of the form

$$p(\theta|\mathcal{I}_{t-1}),\ p(u_{it}|\mathcal{I}_{t-1}),\ t = t_o, \ldots, T, \tag{32}$$

where $\mathcal{I}_{t-1} = \{\mathcal{I}_{i,t-1},\ i = 1, \ldots, n\}$, and t_o represents a starting date that we use to make sequential posterior inferences as data for periods 0, –1, etc., are unavailable. Of course, parameter uncertainty is explicitly accounted for in $p(u_{it}|\mathcal{I}_{t-1}) = p(u_{it}, \theta|\mathcal{I}_{t-1})\, \mathrm{d}\theta$ in (32). The expression can be written further as

$$p(u_{it}|\mathcal{I}_{t-1}) = p(u_{it}|\theta, \mathcal{I}_{t-1}) p(\theta|\mathcal{I}_{t-1})\, \mathrm{d}\theta. \tag{33}$$

If draws $\left\{\theta^{(s)},\ s = 1, \ldots, S\right\}$ are available from $p(\theta|\mathcal{I}_{t-1})$, then the expectation above can be accurately approximated (as $S \to \infty$) by

$$p(u_{it}|\mathcal{I}_{t-1}) = S^{-1} \sum_{s=1}^{S} p(u_{it}|\theta^{(s)}, \mathcal{I}_{t-1}). \tag{34}$$

Of course, the computation is complicated by the fact that we have to condition on the history $u_{t_o:t-1}$ as well which contains latent inefficiencies for all firms from period t_o to $t-1$. Using SMC, it is a simple matter to update the posteriors and approximate accurately the learning elasticity

$$\hat{q}_{\text{active}}(\theta, D) = \frac{\partial \mathbb{E}_{\theta|\mathcal{I}_{t-1}}(u_{it})}{\partial u_{i,t-1}}. \tag{35}$$

Although the active learning approach may provide a more realistic view of learning, it is evident from (35) that it is not quite the same as (31) as the conditioning sets are different. Indeed, (35) conditions on the entire information set $\mathcal{I}_{t-1}$. Although, in principle, this is acceptable, it is more intuitive to condition in merely $\mathcal{I}_{i,t-1}$ or $\mathcal{I}_{i,t-1}$ and a summary of information, denoted $\tilde{I}_{-i,t-1}$ which contains "summary statistics" of all other decision makers (indicated by $-i$). This is not uncommon in the literature and lies at the cornerstone of modern learning theories, see for example Molavi et al. (2018) [70]. In turn, we may be interested in two alternative active learning elasticities. The first is (31) while the second is

$$q_{\text{active},2}(\theta, D) = \frac{\partial \mathbb{E}_{\theta|\mathcal{I}_{i,t-1}, \mathcal{I}_{-i,t-1}}(u_{it})}{\partial u_{i,t-1}}, \tag{36}$$

which conditions on own information and "summary statistics" for the rest of the sector up to date $t-1$. To compute these elasticities, for example the elasticity in (31), we notice that

$$p(u_{it}|\mathcal{I}_{i,t-1}) = \frac{p(u_{it}|\mathcal{I}_{i,t-1})}{p(u_{it}|\mathcal{I}_{t-1})} p(u_{it}|\mathcal{I}_{t-1}), \tag{37}$$

which implies that if we have a draw from $u_{it}|\mathcal{I}_{t-1}$, then this draw is from $u_{it}|\mathcal{I}_{i,t-1}$ with the Metropolis-Hastings acceptance probability $\min\left\{1,\ \frac{p(u_{it}|\mathcal{I}_{i,t-1})}{p(u_{it}|\mathcal{I}_{t-1})}\right\}$. Apparently, the same is the case with $p(u_{it}|\mathcal{I}_{i,t-1}, \mathcal{I}_{-i,t-1})$.
Since $p(u_{it}|\mathcal{I}_{t-1}) = p(u_{it}|\theta, \mathcal{I}_{t-1}) p(\theta|\mathcal{I}_{t-1})\, \mathrm{d}\theta$, if we have MCMC draws for $\theta|\mathcal{I}_{t-1}$we can simulate draws from $u_{it}|\theta, \mathcal{I}_{t-1}$ and, therefore, from $u_{it}|\mathcal{I}_{t-1}$. Therefore, we end up with four different **learning models** (**LM**):

1. The passive learning model in (30).
2. The active learning model in (31).
3. The active learning model in (35).
4. The active learning model in (36).

To implement LM4, we need to decide on the "summary statistics" in $\mathcal{I}_{-i,t-1}$. For each bank, we assume that only the first two moments of inefficiency of all other banks are known to the decision maker i. Using $t_o = 3$ (so that we start in 2004) we compare the different learning models in terms of recursive Bayes factors in panel (a) of

Figure 12: Comparison of different learning models

The benchmark learning model is defined in (19) and (15).

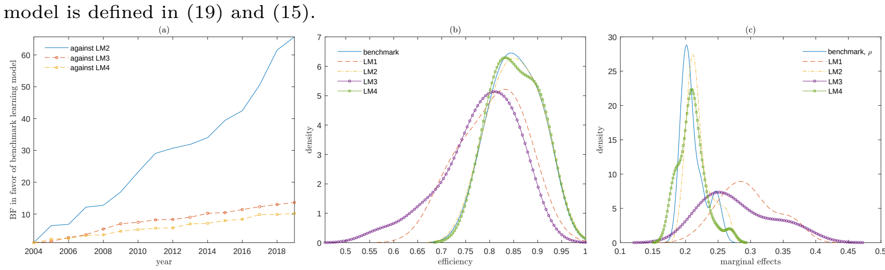

Figure 12. If the benchmark model is indeed "adequate" in the light of the data, its performance should not be too different compared to LM2 and LM4. One piece of evidence is related, of course, to Bayes factors. More evidence is provided by comparing efficiency estimates from LM2 and LM4 with the benchmark (panel (d) of Figure) as well as ρ (panel (e) of Figure 3) with the marginal effects in (31) and (36). Efficiency distributions are compared in panel (b) of Figure 12 and the marginal effects are compared in panel (c).
The evidence in favor of the benchmark learning model (which is defined in (19) and (15)) is decisive against LM3 (the active learning that uses $\mathcal{I}_{t-1}$) but considerably less so for LM2 and LM4.

The fact that LM2 and LM4 are "close" in terms of predictive performance to our benchmark model, implies that the data do not favor a particular model (that conditions on $\mathcal{I}_{i,t-1}$ or $\mathcal{I}_{i,t-1}$ and $\tilde{\mathcal{I}}_{i-,t-1}$) over the benchmark model which is, admittedly, restrictive in terms of its assumptions, compared to the essentially non-parametric formulation in (25).
In panel (b) of Figure we report densities of efficiency estimates for the different learning models. In panel (c) we report densities of the (average) marginal effects with respect to $u_{i,t-1}$. Of course, ρ in (15) is parameter so it does not vary with $u_{i,t-1}$. For the benchmark model, LM2 and LM4 inefficiency distributions are rather close and marginal effects average 0.85 and ranging from 0.67 to nearly 1. The averages of marginal effects are close to 0.22 ranging from 0.15 to less than 0.30. The densities corresponding to LM2 are similar to the benchmark model. For LM1 and LM3 these densities are quite different. For LM1 the marginal effect averages close to 0.3, which is in the right tail of the benchmark model as well as LM2 and LM4, and ranges from 0.15 to about 0.47. For LM3 the marginal effect is somewhat smaller on the average (0.28, median 0.26 and standard deviation 0.051). The standard deviations corresponding to different models are 0.019 (for the benchmark), 0.042 for LM1, 0.019 for LM2, 0.052 for LM3, and 0.021 for LM4. Evidently, they are closer for the benchmark model, LM2, and LM4 but not so for LM1 and LM3.
Overall, there is evidence that the benchmark learning model proposed in this study, LM2 and LM4 behave roughly in the same way albeit this is not the case for LM1 and LM3. LM1 is the passive learning model and LM3 conditions on $\mathcal{I}_{t-1}$, which includes all data up to and including period $t-1$. These differences, as we argued, arise from the fact that the conditioning set matters and it more sensible to condition on a decision maker's information set (plus a few summaries from the other decision makers) rather than rely on marginal effects that assume that all decision makers have full information about the performance of others in all time periods. It is important to point out that ρ in (15) is a single parameter, whereas the marginal effects derived from (25) in their various versions (LM1 through LM4) are much more general constructs that allow for variations of the marginal effect across different values of $u_{i,t-1}$, the conditioning variables, etc. The fact that the posterior density of ρ (conditioning on all the data) and marginal effects from LM2 and LM4 are nearly the same deserves some attention. *First*, the benchmark learning model, LM2 and LM4, despite the great flexibility of the last two, do not behave very differently based on recursive Bayes factors reported in panel (a) of Figure 12. *Second*, in panels (b) and (c) of Figure 12 we consider only average marginal effects in the sake of comparison with (15). Average marginal effects are computed by using the average values of $u_{i,t-1}$ and the other conditional variables but, in general, marginal effects depend on individual values of these variables. Conditioning at other values, could potentially yield different results. We take up this issue in

Figure 13: Comparison of marginal effects for benchmark learning model and LM2, LM4, conditioning on different values of $u_{i,t-1}$ and the conditioning variables.

In panel (a), we compare the marginal effects from the benchmark learning model (heavy line) and LM2. In panel (b) we compare the marginal effects from the benchmark learning model and LM4. For LM2 and LM4 we consider 50 randomly chosen values from the support of $u_{i,t-1}$ and the conditioning variables.

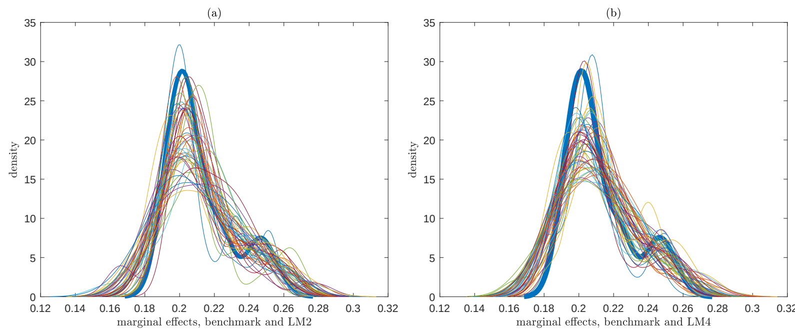

Figure 13, where we examine the effect of conditioning on other values.
In panel (a) of Figure 13, we compare the marginal effects from the benchmark learning model (heavy line) and LM2. In panel (b) we compare the marginal effects from the benchmark learning model and LM4. For LM2 and LM4 we consider 50 randomly chosen values from the support of $u_{i,t-1}$ and the conditioning variables.

Clearly, conditioning on different values of the predictors in (25), provides distributions of marginal effects that are different, although the differences are not so much in terms of location and spread. For example, bimodality of the posterior density of ρ (near the modes 0.20 and 0.25) is shared by many of the distributions of marginal effects when we condition not on the average but other values of the predictors in (25).
Based on the nonlinear representation in (25) we can also derive long-run or steady-state inefficiency. The sample distributions of steady-state inefficiency for LM2 and LM4, are presented in panel (a) of Figure 14 where we also present the steady-state distribution for the benchmark model. The steady-state densities turn out to be bimodal having a major mode near 0.18 and another near 0.25. Average steady-state inefficiency is 0.186 with standard deviation close to 0.030 (the medians of LM2 and LM4 are 0.176). So, we expect that, in the long-run, inefficiency averages close to 18%. In panel (b) we report results for LM1 and LM3.

Results from LM1 and LM3, reported in panel (b), are quite different. For example, LM1 has steady-state inefficiency which averages 0.124 (0.119 for LM3) with standard deviation 0.027 (0.015 for LM3). Therefore, the implications of (LM2, LM4) on the one hand, and (LM1, LM3) are quite different in terms of steady-state inefficiency; the second group, for example implies much lower steady-state inefficiencies and different distributions relative to the first group. The distributions of steady-state inefficiency in models LM1 and LM3 are also not the same.

7 Can we attribute all dynamics to learning?

7.1 General

An important question when we compare the benchmark learning model to models in (25) is whether all dynamics in (25) and, therefore, (30), (31), (35), and (36), can be attributed to learning. The fact that the benchmark, LM2 and LM4 provide, roughly, the same implications is, of course, quite re-assuring in this direction. Additional

Figure 14: Steady-state inefficiency densities

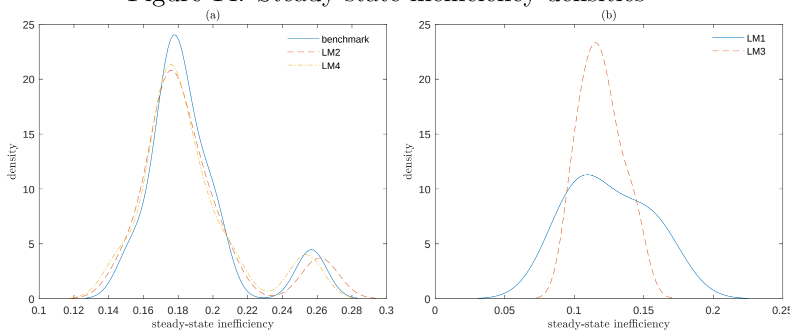

evidence can be provided if we decompose inefficiency as

$$u_{it} = \overline{u}_{it} + \widetilde{u}_{it}, \tag{38}$$

where $\overline{u}_{it}$ represents the non-learnable component and $\widetilde{u}_{it}$ denotes the component that can be learned. Therefore, we split efficiency in two parts, the first of which may be dynamic but is not part of learning (it is, for example, persistent inefficiency), and the second is the component that managers and decision makers can actually learn. Total efficiency, or u_{it}, can be obtained from LM2 in (31) or LM4 in (36) as these models were found to be, in many respects, similar to the benchmark. The $\widetilde{u}_{it}$ component, in principle, follows (9) whereas $\overline{u}_{it}$ is obtained residually.[8] If we proceed along this direction, we obtain estimates of $\overline{u}_{it}$ or the ratio

$$\psi_{it} \equiv \frac{\overline{u}_{it}}{u_{it}},$$

which represents the proportion of non-learnable inefficiency. The sample distributions of ψ_{it} are reported in Figure 15.

The evidence reported in Figure 15, suggests that *the portion of inefficiency that cannot be learned is, in fact,* sizable, and averages 30% in both LM2 and LM4 (standard deviations close to 3% and 2.8%, respectively, and values ranging between 20% and 49%). In turn, this suggests that in U.S. banking a large part of inefficiency (close to 70%) can be learned and reduced. The fact that it is not reduced can be attributed to a number of reasons, the most prominent of which is the "quiet life hypothesis", first suggested by Hicks. According to this hypothesis, banks with higher market power can have high profitability even though it could cause inefficiency. This can be expressed equivalently by saying that managers in a non-competitive environment will avoid risky or difficult decisions, and do not have the initiative or incentives to engage in projects that reduce waste and inefficiency (Koetter and Vins, 2008 [64]; Koetter et al., 2012 [63]). More importantly, perhaps, the fact that ψ_{it} is "small", shows that, at least for the most part, (31) and (36), are consistent with *learning* dynamics (as opposed to general inefficiency dynamics) and, specifically, the dynamics consistent with our benchmark model in (9), (8), and (15).

7.2 What can we learn from the non-learnable component?

The components $\overline{u}_{it}$ and $\widetilde{u}_{it}$ are, potentially, important in establishing more properties of the learning process. First, an interesting question is whether the learnable component has, roughly, the same properties as in Section 4 and Figures (3)-(6). Second, from the non-learnable component we would like to know what exactly is it about it that prevents learning. As to the first question, the empirical evidence (not reported here but available on request) shows that the learnable component has the same properties as in Section 4 and Figures (3)-(6) where the presence of $\overline{u}_{it}$ has been ignored. This is surprising, as (15) shows that this component is not trivial. To understand the

[8] We follow the so-called corrected least squares approach, and use the estimates $\widetilde{u}_{it} - \min_{i,t} \widetilde{u}_{it}$ to satisfy the non-negativity restriction.

Figure 15: Sample densities of ψ_{it}

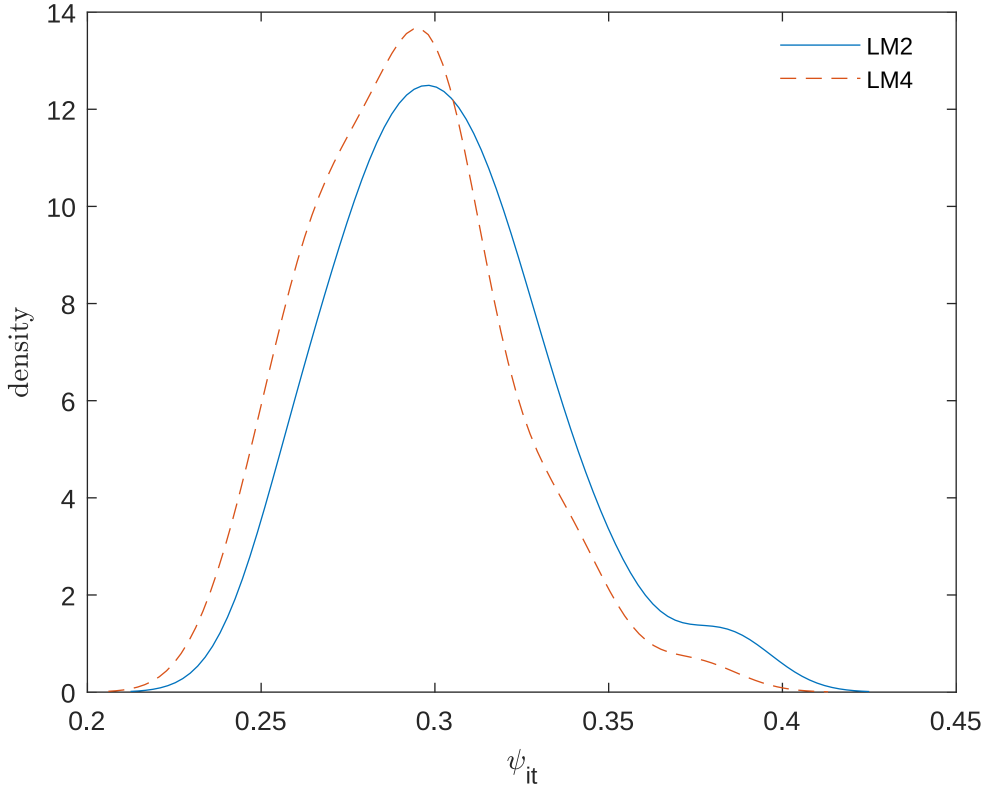

Figure 16: Marginal effects of the non-learnable component, Part I

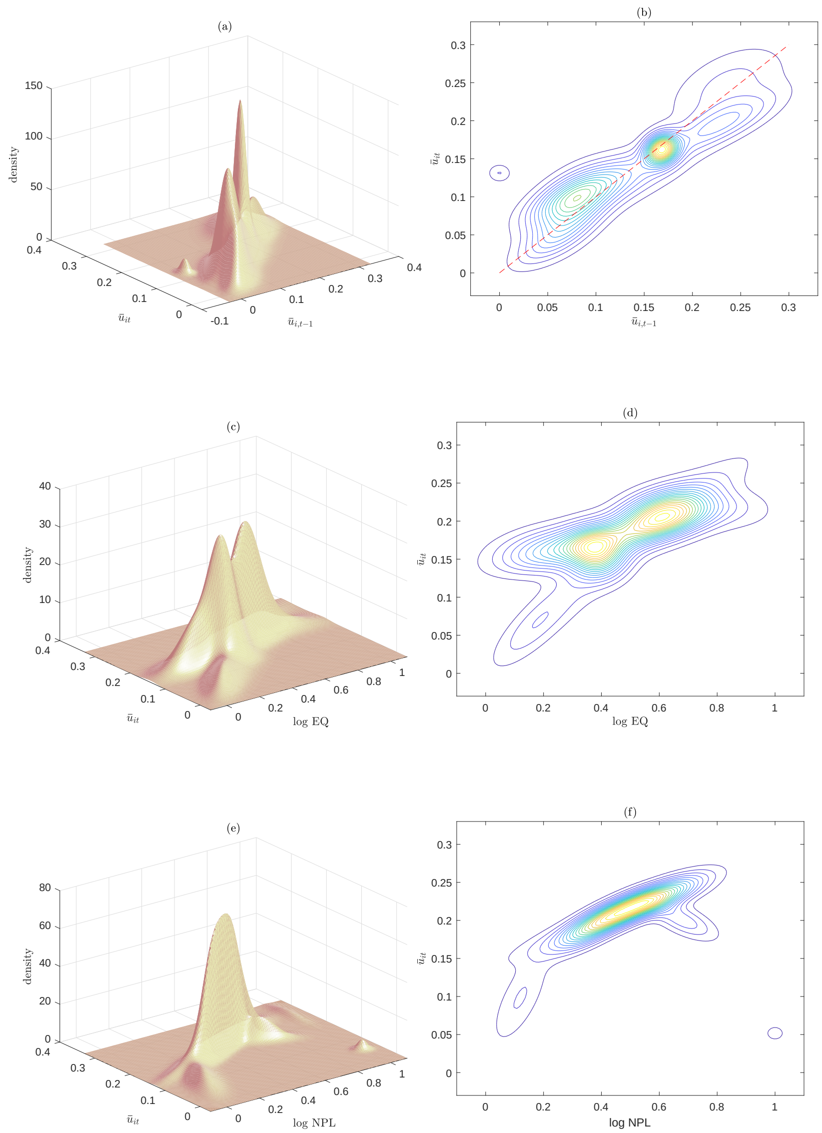

properties of $\overline{u}_{it}$ we consider the following model

$$\overline{u}_{it} = \varphi\left(\overline{u}_{i,t-1}, \log EQ_{i,t-1}, \log NPL_{i,t-1}, \mathscr{P}_{i,t-1}, \mathscr{T}_{it}\right) + \xi_{it}^{u} + \eta_{i}^{u}, \tag{39}$$

where $\varphi(\cdot)$ is an unknown functional form, $\overline{u}_{i,t-1}$ represents the lagged value of the non-learnable component, $\log EQ_{i,t-1}$ is the lagged log value of equity, $\log NPL_{i,t-1}$ is the lagged log value of non-performing loans, $\mathscr{P}_{i,t-1}$ represents the lagged value of the first principal component of inputs and outputs (separately), [9], $\mathscr{T}_{it}$ represents the time trend, ξ_{it}^{u} represents an error term, and η_{i}^{u} represent bank-specific effects, respectively. If the non-learnable component does not depend on the specific set of variables then it can be interpreted as a non-causal part of performance that cannot be learned objectively.

To estimate (39) we use an artificial neural network whose order is chosen by the data using the marginal likelihood criterion. The details are explained in Appendix C.

In panels (a) and (b) of Figure 16, we present the distribution of marginal effects $\frac{\partial \varphi(\cdot)}{\partial u_{i,t-1}}$ which appears to be

[9] The first principal components account for almost 80% of the total variation in logs of inputs and outputs.

Figure 17: Marginal effects of the non-learnable component, Part II

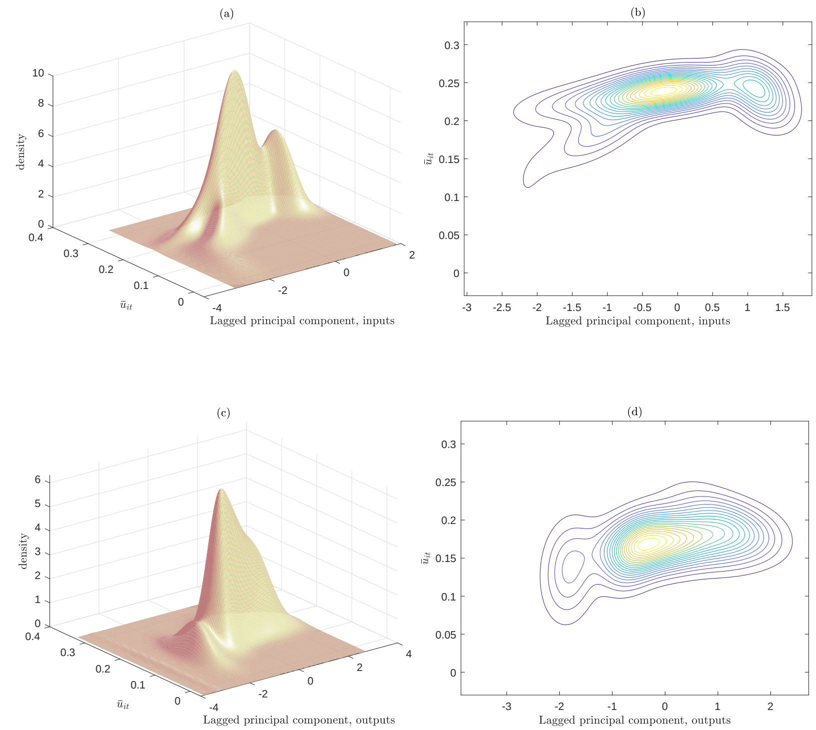

positive but multimodal (also shown on panel (b) is the 45° degrees line). This evidence shows that, in fact, $\overline{u}_{it}$, is persistent so, this component is non-learnable not because it is random but because it is genuinely non-learnable. As a function of equity and non-performing loans, the respective marginal effects are shown in panels (c), (d), and (e) and (f), respectively.

As a function of the first principal components, the association is weaker but still, evidently, positive; see panels (a)-(d) of Figure 17. The message that emerges from this evidence is that *the non-learnable component is highly persistent, depends positively on equity and NPLs, and it is not independent of inputs and outputs*. High persistence indicates that, ceteris paribus, $\overline{u}_{it}$, is, indeed, difficult to learn, particularly so in larger banks (notice the effect of equity), with more NPLs (notice the positive relation with NPLs). The size effect is confirmed by the association between inputs and outputs in Figure 17. Of course, banks with higher NPLs also show a size effect but NPLs also reflect the effect of uncertainty so, in larger organizations, operating in more uncertain environments, the relationships between the non-learnable component of performance (inefficiency) are highly complex (evidenced by the multimodalities of the joint distributions in Figures 16 and 17), depend, on the main, positively on size and uncertainty and they are heterogeneous (in the sense that multimodalities are compatible with different clusters or groups in the data). As noted in Arbelo, Arbelo-Pérez and Pérez-Gómez (2021), “the relationship between efficiency and size heavily depends on the internal properties and characteristics of the firm and environment in which it operates and that there is heterogeneity among firms”
It is known from the neuroscience literature that, at a fundamental level, learning is about forming, removing, or changing associations (Bassett and Mattar, 2017). Provided these associations are close to the ones predicted by the neural network results in Figures 16 and 17 (Gadzinski and Castello, 2020), it turns out that the non-learnable component has extensive grounding in the organizational structure as it depends on its observable outcomes. Although this finding is in line with major points in the literature that learning depends on past experience (Minniti and Bygrave, 2001), *it is also a new contribution in the sense that the uncertainty in the environment can be quantified more precisely and analyzed in a systematic manner*, as in (39).[10]

10

Gadzinski and Castello (2020) decompose uncertainty into model-based and error-based. The first is due to deficiencies of the benchmark model and the second is unavoidable. They assume the two errors are independent although this is not always the case in

As the non-learnable performance component depends systematically on observable variables, one may be tempted to use the associations provided by (39) and presented in Figures 16 and 17, to "learn" about the "non-learnable" component or at least devise a policy that reduces its magnitude. Although this may be an interesting avenue for future research, it must be emphasized that (i) one has to verify that the associations implied by (39) are causal, and (ii) given that they are causal, develop a model of how the organizational structure is reflected at the macro level associations provided by (39). This reflection will, undoubtedly, be compatible with multiple narratives or theories about how underlying structures are reflected in performance; and, clearly, how such reflections are moderated or mediated by other non-observables or practices in the organization.

Concluding remarks

In this paper we have proposed and implemented a model of Bayesian learning for technical inefficiency in stochastic frontier models. The model is estimated using Bayesian techniques organized around Markov Chain Monte Carlo using data on large U.S. banks for the period 2000-2019. The evidence suggests that there is some learning in terms of reducing inefficiency, although this effect is more pronounced during the subprime crisis. Jump sizes and probabilities do not contain autoregressive elements so, they are not persistent, and for the entire sample there seems to be no relationship between productivity growth, jump sizes and jump probabilities. During the subprime crisis, however, the relationship is negative and much more pronounced showing that lower productivity growth increases the probability of a jump and the jump size in experience itself.

We validate the model (which we call "benchmark learning model") using a non-parametric neural network model that allows arbitrary learning dynamics. Close to the average values of the predictors, dynamic marginal effects are quite close suggesting that the proposed benchmark model is a good representation of the data. Globally, recursive Bayes factors suggest that the benchmark and the non-parametric neural network model have similar predictive ability. In the non-parametric neural network model we distinguish between three possible types of learning. First, passive learning, in which all data is used to make statistical inferences about performance learning. Second, active learning, where the decision maker learns based on his / her own experience. Third, when the decision maker is allowed to learn based, additionally, on summary statistics from the performance distribution of other decision makers. Critical in the performance of the benchmark learning model is the comparison with these different versions that condition on different information sets. Passive learning is shown not to perform as well as the benchmark, whereas versions of active learning models perform equally well and deliver results that are, qualitatively and quantitatively, quite similar to the benchmark, despite the fact that the benchmark is simple.

In further investigations, we decompose performance into a component that can be learned and potentially reduced and another part that can be, naturally, attributed to complexities of the environment, the realities of organizational structure, and how these are reflected at the non-learnable component and its underlying fundamental. The underlying fundamental are shown to be related to observables such as equity, non-performing loans, inputs, and outputs so, the non-learnable component reflects more the underlying organizational structure rather than components which are beyond the reach of learning. These components can, in fact, be learned, as they are associated with fundamentals in the organizational structure and observed variables at the macro level. Similarly, a large part of overall inefficiency (close to 70%) can be learned and reduced. In the second case, the fact that it is not reduced means that it is not efficient to reduce it (e.g. "quiet life hypothesis"). In the first case, it means that investing time and resources to modify organizational structures, is not worth the effort. Therefore, our model, first, provides a picture of performance and what type of learning is applied to improve it; and, second, it decomposes overall performance into a component that can be learned (and reduced, if desired) and a component that reflects the complexities of the environment and the organizational structure. Taking into account that average inefficiency is close to 15% (panel (b) of (12)), the extent of non-learnable inefficiency is close to 4.5 percentage points, 15% times the average 30% of the non-learnable component). This translates to, roughly, more than 45 million USD per bank in constant 2005 prices. For the entire banking sector over the period 2000-2019, this cost translates to a lower bound of 256,5 billion USD in constant 2005 prices. The inefficiency that can, in fact, be learned and reduced, has a lower bound of 589.8 billion USD, giving a rough total of 885 billion. As several banks have total assets in excess of a billion, this number is only a lower bound showing that there are incentives, for the regulatory authorities at least, to restrict waste in the sector.

The finding that performance is, at least in the main, subject to learning and this finding is robust and reliable, is

practice if relevant explanatory variables have been omitted from the model becoming part of the error. As out focus here is performance, part of u_{it}, viz. $\widetilde{u}_{it}$ can be learned but another part, $\overline{u}_{it}$, cannot. The second component is part of the complexities characterizing the environments provided it does not contain specification errors. We use the standard econometric diagnostics for autocorrelation, heteroskedasticity, nonlinearity (RESET test) in the residuals produced from (39) at the posterior means of the parameters.

important on its own and it is consistent with alternative theories, for example, the Hicksian "quiet life hypothesis" (revealing the importance of market power in the U.S. banking sector). To our knowledge, the decomposition and estimation of a "non-learnable" component is novel and so is the finding that it averages 30%, ranging, however, between 20% and, nearly, 42% of total inefficiency in the U.S. banking sector. Similar decompositions can be applied to other data sets which can, hopefully, lead to a better understanding of learning in organizations. The non-learnable component, in the present study, is found to be highly persistent and associated with observed variables such as equity, non-performing loans, as well as inputs and outputs. From this point of view, it is not epistemically non-learnable but rather it reflects complexities and realities of the organizational structure that are persistent and hard to change. The observed associations may well be non-causal, that is they may change when the environment changes but this, again, reflects the deep complexities of what is the "environment" or the "organizational structure" in themselves. Broad definitions of environment relate to "situations" or constellations consisting of opportunities, relations, processes, etc.[11]

Although learnable and non-learnable performance components share a common environment, if we can document, as in this study, that there is robust and reliable learning model, confined to appropriate information sets, then, residually, the non-learnable component should not contain, at least *in principle*, no further information about the learning process. The problem is that this "principle" although it may be true as a statistical construct (which is predicated on orthogonality or independence between errors and the main components of the learning model) it may fail to be true once we condition on more realistic information sets that reflect the actual learning capabilities of organizations. In this study, we have shown that this conditioning makes a difference and results from different conditionings or information sets, can not only be compared but also evaluated as in regular comparison of statistical hypotheses. Conditioning in different information sets is operationalized using standard tools from the arsenal of Sequential Monte Carlo, opening the way to more applications in diverse organizational domains.

Technical Appendix

A.1 Time trend model (Model I)

Bayesian analysis is facilitated using Markov Chain Monte Carlo (MCMC, Tsionas, 2000 [87]) to draw from the posterior conditional distribution of $u_{it}|\theta, D$. We draw from the posterior conditional distribution of β and σ_v as follows.

$$\beta|\mathbf{u}, \sigma_v, \lambda, \sigma_\theta, D \sim \mathcal{N}_K\left(b, V\right), \tag{A.1}$$

$$\tfrac{(\mathbf{y}+\mathbf{u}-\mathbf{X}\beta)'(\mathbf{y}+\mathbf{u}-\mathbf{X}\beta)}{\sigma_v^2}|\beta, \lambda, \sigma_\theta, \mathbf{u}, D \sim \chi^2_{nT}, \tag{A.2}$$

where $b = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'(\mathbf{y}+\mathbf{u})$, $V = \sigma_v^2(\mathbf{X}'\mathbf{X})^{-1}$ and $\mathbf{y}, \mathbf{X}$ denote, respectively, the $nT \times 1$ and $nT \times K$ arrays of all observations on the dependent variable and the regressors. The joint conditional posterior of λ, σ_θ is given by the following expression.

$$p(\lambda, \sigma_\theta|\beta, \mathbf{u}, D) \propto \lambda^{-(nT/2+1)}\sigma_\theta^{-(nT+1)}\prod_{i=1}^{n}\prod_{t=1}^{T}\left[(\lambda+\tau_{it})^{1/2}\exp\left\{-\tfrac{u_{it}(\lambda+\tau_{it})}{2\lambda\sigma_\theta^2}\right\}\right]. \tag{A.3}$$

The conditional posterior of σ_θ is inverted *gamma*, viz. $\frac{\sum_{i=1}^{n}\sum_{t=1}^{T}(\lambda+\tau_{it})}{\lambda\sigma_\theta^2}|\beta, \mathbf{u}, \lambda, D \sim \chi^2_{nT}$. The conditional posterior of λ is not in a known family but we can use rejection sampling from an inverted *gamma* distribution of the form $q(\lambda;\alpha) = \frac{(\alpha/2)^{nT}}{\Gamma(nT)}\lambda^{-(nT/2+1)}e^{-(\alpha/2)\lambda}$ where the parameter α is determined to maximize the acceptance rate. The optimal choice is $\alpha = \frac{nT}{\hat{\lambda}}$ and it is the value the solves the saddle point problem: $\underset{\alpha}{\min}\underset{\lambda}{\max}\frac{p(\lambda|\beta,\sigma_\theta,\mathbf{u},D)}{q(\lambda;\alpha)}$; where $\hat{\lambda}$

[11] "[A] situation defines an abstract state of affairs that represents a particular scenario of interest and can consist of resources, relations and processes. In other words, a situation is a combination of one or several resources, and / or one or sever allocations, or one or more sensor measurements linked through spatial, temporal and / or spatio-temporal relationships." (Schreiber et al., 2000; see also Giustozzi et al., 2018, p. 683).

solves the nonlinear equation

$$\sum_{i=1}^{n}\sum_{t=1}^{T}\left(\hat{\lambda}+\tau_{it}\right)^{-1}+\frac{nT}{\hat{\lambda}}=\frac{\sum_{i=1}^{n}\sum_{t=1}^{T}u_{it}\tau_{it}}{\sigma_{\theta}^{2}}. \tag{A.4}$$

A.2 Particle Filtering (Models II and III)

We use the sequential PF of Chopin (2002). Given a target posterior $p(\theta|Y) := p(\theta|Y_{1:T})$ a particle system is a sequence $\{\theta_j, w_j\}$ such that $E(h(\theta)|Y) := h(\theta)p(\theta|Y)d\theta \cong \lim_{J\to\infty} \frac{\sum_{j=1}^{J} w_j h(\theta_j)}{\sum_{j=1}^{J} w_j}$, almost surely, for any measurable function h, provided the expectation exists. We consider the sequence of posterior distributions $p_t := p(\theta|Y_t)$. The PF algorithm is as follows.
Step 1. Reweight: update the weights $w_j \leftarrow w_j \frac{p_{t+1}(\theta_j)}{p_t(\theta_j)}, j = 1, \ldots, J$.
Step 2: Resampling: resample $\{\theta_j, w_j J_{j=1}^{H} \to \{\theta_j^r, 1\}_{j=1}^{J}$.
Step 3. Move: draw $\theta_j^m \sim K_{t+1}(\theta_j^r), j = 1, \ldots, J$, where K_{t+1} is any transition kernel whose stationary distribution is p_{t+1}.
Step 4. Loop: $t \leftarrow t+1, \{\theta_j, w_j\}_{j=1}^{J} \leftarrow \{\theta_j^m, 1\}_{j=1}^{J}$ and return to Step 1.
Chopin (2002) [66] recommends the independence Metropolis algorithm to select the kernel, which requires a source distribution. A possible choice, if we sampled from p_n $(n < T)$, with respect to p_{n+s} is $\mathcal{N}(\hat{E}_{n+s}, \hat{V}_{n+s})$ where

$$\hat{E}_{n+s} = \frac{\sum_{j=1}^{J} w_j \theta_j}{\sum_{j=1}^{J} w_j},\ \hat{V}_{n+s} = \frac{\sum_{j=1}^{J} w_j \left(\theta_j - E_{n+p}\right)\left(\theta_j - E_{n+p}\right)'}{\sum_{j=1}^{J} w_j}.$$

The strategy can be parallelized easily. If K processors are available, we can partition the particle system into K subsets, say $S_k, k = 1, \ldots, K$), and implement computations for particles of S_k in processor k. The algorithm can deal with new data at a nearly geometric rate and therefore the frequency of exchanging information between processors (after reweighting) decreases at a rate exponential to n, which is highly efficient.
Resampling according to $\theta_j^m \sim K_t(\theta_j^r, .)$ reduces particle degeneracy (Gilks and Berzuini, 2001 [38]) since identical replicates of a single particle are replaced by new ones without altering the stationary distribution. For this application using $J = 2^{12}$ particles gave a mean squared error in posterior means of 10^{-5} over 100 runs.
Chopin (2004) [22] introduces a variation of MSC in which the observation dates at which each cycle terminates (say $t_1, \ldots, t_L$) and the parameters involved in specifying the Metropolis updates (say $\lambda_1, \ldots, \lambda_L$) are specified. Therefore, $0 = t_0 < t_1 < \cdots < t_L = T$ and we have the following scheme (we rely heavily on Durham and Geweke, 2013 [29]).
Step 1. Initialize $l = 0$ and $\theta_{jn}^{(l)} \sim p(\theta)$, $j \in \mathcal{J}, n \in \mathcal{N}$.
Step 2. For $l = 1, \ldots, L$:
(a) Correction phase:
(i) $w_{jn}(t_{l-1}) = 1$, $j \in \mathcal{J}, n \in \mathcal{N}$
(ii) For $s = t_{l-1} + 1, \ldots, t_l$

$$w_{jn}(s) = w_{jn}(s-1)p(y_s|y_{1:s-1}, \theta_{jn}^{(l-1)}),\ j \in \mathcal{J}, n \in \mathcal{N}.$$

(iii) $w_{jn}^{(l-1)} := w_{jn}(t_l)$, $j \in \mathcal{J}, n \in \mathcal{N}$.
(b) Selection phase, applied independently to each group $j \in \mathcal{J}$: Using multinomial or residual sampling based on $\left\{w_{jn}^{(l)}, n \in \mathcal{N}\right\}$, select

$$\{\theta_{jn}^{(l,0)}, n \in N\}$$

from $\{\theta_{jn}^{(l-1)}, n \in N\}$.
(c) Mutation phase, applied independently across $j \in \mathcal{J}, n \in \mathcal{N}$:

$$\theta_{jn}^{(l)} \sim p(\theta|y_{1:t}, \theta_{jn}^{(0)}, \lambda_l) \tag{A.5}$$

Figure 18: Prior sensitivity analysis

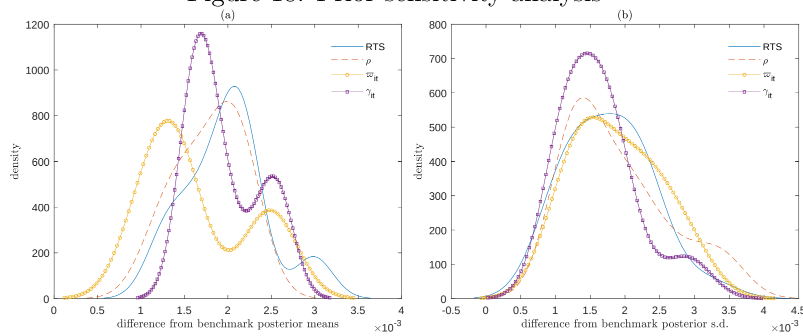

where the drawings are independent and the pdf above satisfies the invariance condition:

$$ {}_{\Theta}\, p(\theta|y_{1:t_l},\theta^*,\lambda_l)p(\theta^*|y_{1:t_l})d\nu(\theta^*) = p(\theta|y_{1:t_l}). \tag{A.6} $$

Step 3. $\theta_{jn} := \theta_{jn}^{(l)}$, $j \in \mathcal{J}, n \in \mathcal{N}$.
At the end of every cycle, the particles $\theta_{jn}^{(l)}$ have the same distribution $p(\theta|y_{1:t_l})$. The amount of dependence within each group depends upon the success of the Mutation phase which avoids degeneracy.

A.3 Prior sensitivity analysis and numerical performance of simulation techniques

The critical prior parameters in our analysis are $\bar{s}, \bar{q}, \bar{n}, \bar{N}, \bar{Q}$ defined in section 3. We select 1,000 randomly chosen values for these parameters. For $\bar{s}, \bar{q}, \bar{Q}$ we generate values using a uniform distribution in the interval $\left(10^{-5}, 10\right)$ while for $\bar{n}$ and $\bar{N}$ we use a uniform distribution in the interval (1, 10). We use the Gibbs sampler for Model I and the PF algorithm for Models II and III. To minimize computational burden, we use 15,000 iterations (starting from the benchmark posterior means) omitting the first 5,000 in the burn-in period. Across the 1,000 different priors we would like to see whether key aspects of the model like returns to scale, productivity growth, ρ, ϖ_{it} and γ_{it} are sensitive to the prior.
MCMC can be seriously affected if the draws are highly autocorrelated. Of course, some autocorrelation is unavoidable but the empirical issue is whether it is so large as to prevent thorough exploration of the posterior. Autocorrelation can be detected using the autocorrelation functions (acf) corresponding to MCMC draws. Additionally, one can look into relative numerical efficiency (RNE, Geweke, 1992 [35]) which is equal to one in the case of i.i.d draws from the posterior. Finally, convergence can be tested using Geweke's (1992) [35] z-diagnostic known as Geweke Convergence Diagnostic (GCD).
As we have a large number of parameters, we report maximum autocorrelations across all parameters, minimum RNE and maximum GCDs in Figure A.2.
Our results are reported in Figure A.1 and show differences of the posterior means corresponding to the new priors relative to the posterior means obtained through the benchmark prior in Section 3. Evidently, the difference in posterior means (panel(a)) and posterior standard deviations (panel (b)) are fairly small so, the impact of the prior information is not excessive.

The acf's are reported in panel (a) of Figure A.2, RNE densities in panel (b), and GCDs in panel (c). From this evidence, although there is autocorrelation, this does not affect MCMC convergence (panel (c)) and, certainly, autocorrelation is not destructively large.

Figure 19: Simulation performance diagnostics

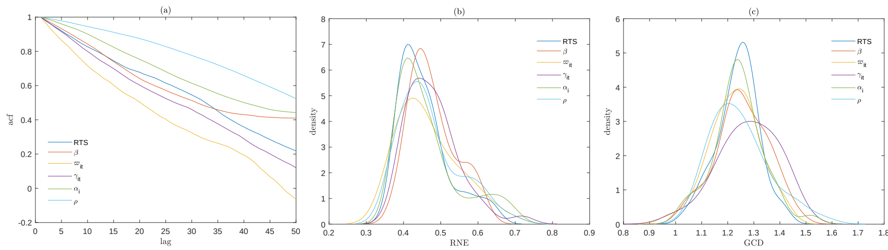

Notes: "RTS" denotes returns to scale, and α_is are bank effects in (15).

Appendix C

In this Appendix, we describe our decomposition of (39) into learnable and non-learnable components. The obvious decomposition is to identify the learnable component with the fitted values of the model, and the non-learnable component with the residuals from estimating (39). Although this decomposition is time honored (Barro, 1977, 1978, 1979; Lucas, 1973), we believe it is not appropriate because the residuals condition on the entire data set and, therefore, (i) this does not allow a sequential interpretation which is closer to the realities of learning, and (ii) it does not account for the fact that learning conditions on the realities and structures of a particular organization, operating in a certain environment or "situation".

We write (39) in a more general form as follows:

$$Y_{it} = \varphi(X_{it}; \beta) + \eta_i + e_{it}, \tag{C.1}$$

where Y_{it} is the dependent variable (inefficiency, in our case), X_{it} is a vector of regressors, $\varphi(\cdot; \beta)$ represents (39), β is a parameter vector, η_i represents organization-specific effects, and e_{it} is a statistical error term.

First, one should ensure that (C.1) is statistically acceptable in the sense that it passes certain diagnostic tests for heteroskedasticity, autocorrelation, etc. Otherwise, specification errors become part of the error term and, as a result, error decompositions become ambiguous.

Second, if the specification in (C.1) is statistically acceptable, we produce estimates of the learnable and non-learnable components (that are organizational specific) as follows. The model provides an overall posterior distribution $p(\beta, \eta|D)$ where D is the entire data set on $\{X_{it}, Y_{it}\}$. As a result, we can define the model errors

$$e_{it} = Y_{it} - \varphi(X_{it}; \beta) - \eta_i, \tag{C.2}$$

where η is the vector of all individual effects (and possibly other nuisance parameters). Model errors in (C.2) are no different compared to (C.1) so, (C.2) is not operational in itself. It becomes operational, once we define

$$\hat{e}_{it}^{(s)} = Y_{it} - \varphi(X_{it}; \beta^{(s)}) - \eta_i^{(s)},\ 1 \le s \le S, \tag{C.3}$$

where $\left\{\beta^{(s)}, \eta^{(s)}\right\}$ represent MCMC draws from the posterior $p(\beta, \eta|D)$. Again, Bayesian residuals (Chaloner and Brant, 1988) in (C.3) condition on the entire data set and, therefore, they do not reflect the realities of learning. Rather, we are interested in the posterior

$$p(\beta_i, \eta_i | D_{i,t_o:t-1}), \tag{C.4}$$

where $D_{i,t:t-1}$ denotes data for bank i for periods t_o through $t-1$ (where t_o is the initial time period that we set as in the discussion of (30)-(36)), β_i, η_i are bank-specific parameters, and $p(\beta_i, \eta_i | D_{i,t_o:t-1})$ denotes the posterior of parameters for the ith bank conditional on information of this bank from period t_o through $t-1$. Suppose we have a set of MCMC draws $\left\{\beta_i^{(s)}, \eta_i^{(s)},\ 1 \le s \le S\right\}$ from (C.4), then we can define the new residuals

$$\tilde{e}_{it}^{(s)} = Y_{it} - \varphi(X_{it}; \beta_i^{(s)}) - \eta_i^{(s)},\ 1 \le s \le S. \tag{C.5}$$

The constructions in (C.3) and (C.5) are quite different as (i) they condition on different information sets, and (ii) they allow for heterogeneity in the form of bank-specific parameters that are obtained sequentially for each time period in the sample. Our performance measures in (39) are related to these constructions as follows.

$$\begin{aligned} \overline{u}_{it}^{(s)} &:= e_{it}^{(s)},\ 1 \le s \le S, \\ \widetilde{u}_{it}^{(s)} &:= \varphi(X_{it}; \beta_i^{(s)}) - \eta_i^{(s)},\ 1 \le s \le S. \end{aligned} \tag{C.6}$$

The only remaining problem is to obtain MCMC draws from (C.4) provided we have MCMC draws $\left\{\beta^{(s)}, \eta^{(s)},\ 1 \le s \le S\right\}$ and, therefore, access to the posterior $p(\beta, \eta | D)$. This can be accomplished using the overall MCMC draws $\left\{\beta^{(s)}, \eta^{(s)},\ 1 \le s \le S\right\}$ by accepting the candidate $\beta^{(s)}, \eta^{(s)}$ with the Metropolis-Hastings probability

$$\min\left\{1,\ \frac{p(\beta^{(s)}, \eta^{(s)} | D_{i,t_o:t-1})}{p(\beta^{(s)}, \eta^{(s)} | D)}\right\}. \tag{C.7}$$

References

[1] Ambad, S.N.A., & Damit, D.H.D.A., 2016. Determinants of entrepreneurial intention among undergraduate students in Malaysia. Procedia Econ. Finan. 37, 108–114.

[2] Anzanello, M.J., Fogliatto, F.S., & Santos, L., 2014. Learning dependent job scheduling in mass customized scenarios considering ergonomic factors. Int. J. Prod. Econ. 154, 136–145.

[3] Arbelo, A., Arbelo-Pérez, M., & Pérez-Gómez, P. (2021). Are SMEs less efficient? A Bayesian approach to addressing heterogeneity across firms. Small Business Economics, in press.

[4] Arrow, K. J. (1962). The economic implications of learning-by-doing. Review of Economic Studies, 29: 155–173.

[5] Arrow, K. J. (2014). Commentary. In J. A. Scheinkman (ed.) Speculation, Trading, and Bubbles, pp. 81–88. New York: Columbia University Press.

[6] Arrow, K. J. (2016). Commentary. In J. E. Stiglitz, B. C. Greenwald (eds) Creating a Learning Society: A New Approach to Growth, Development, and Social Progress, pp. 504–508. New York: Columbia University Press.

[7] Azadegan, A., Srinivasan, R., Blome, C., & Tajeddin, K., 2019, Learning from near-miss events: An organizational learning perspective on supply chain disruption response. International Journal of Production Economics 216, 215–226.

[8] Badunenko, O., & S. C. Kumbhakar, 2017. Economies of scale, technical change and persistent and time-varying cost efficiency in Indian banking: Do ownership, regulation and heterogeneity matter? European Journal of Operational Research 260 (2), 789–803.

[9] Barba-Sanchez, V., & Atienza-Sahuquillo, C., 2018. Entrepreneurial intention among engineering students: the role of entrepreneurship education. European Res. Manag. Bus. Econ. 24 (1), 53–61.

[10] Barro, R. J. (1977), Unanticipated Money Growth and Employment in the United States. American Economic Review 67, 101–115.

[11] Barro, R. J. (1978), Unanticipated Money Growth, Output and the Price Level in the United States. Journal of Political Economy 86 549–580.

[12] Barro, R. J. (1979). Unanticipated Money Growth and Unemployment in the United States: Reply. American Economic Review 69, 1004–1009.

[13] Bassett, D. S., & Mattar, M. G., 2017. A Network Neuroscience of Human Learning: Potential to Inform Quantitative Theories of Brain and Behavior. Trends in Cognitive Sciences 21 (4), 250–264.

[14] Bogetoft, P., & J. L. Hougaard, 2003. Rational inefficiencies. Journal of Productivity Analysis, 20, 243–271.

[15] Boussemart, J-P., Leleu, H., Shen, Z., Vardanyan, M., & Zhu, N. (2019). Decomposing banking performance into economic and credit risk efficiencies. European Journal of Operational Research 277 (1), 719–726.

[16] Camanho, A. S., & R. G. Dyson, 2005. Cost efficiency measurement with price uncertainty: a DEA application to bank branch assessments. European Journal of Operational Research 161 (2), 432–446.

[17] Camerer C. (1998). Bounded Rationality in Individual Decision Making. Experimental Economics 1, 163–183.

[18] Casu, B., A. Ferrari, C. Girardone, & J. O. S. Wilson, 2016. Integration, productivity and technological spillovers: Evidence for eurozone banking industries. European Journal of Operational Research 255 (3), 971–983.

[19] Cavagnini, R., Hewitt, M., & Maggioni, F., 2020. Workforce production planning under uncertain learning rates. Int. J. Prod. Econ. 225, 107590.

[20] Chaloner, K., & Brant, R., 1988, A Bayesian approach to outlier detection and residual analysis. Biometrika, 75, 651–59.

[21] Chopin, N. (2002). A sequential particle filter method for static models, Biometrika 89 (3) 539–551.

[22] Chopin, N. (2004). Central limit theorem for sequential Monte Carlo methods and its application to Bayesian inference. Annals of Statistics 32, 2385–2411.

[23] Deakins, D., & M. Freel (1998). Entrepreneurial learning and the growth process in SMEs. The Learning Organization 5 (3), 144–155.

[24] Delis, M., M. Iosifidi, & M. G. Tsionas, 2017. Endogenous bank risk and efficiency. European Journal of Operational Research, 260 (1), 376–387.

[25] Delis, M., M. Iosifidi, & M. G. Tsionas, 2020. Management estimation in banking. European Journal of Operational Research 284 (1), 355–372.

[26] DiCiccio, T. J., Kass, R. E., Raftery, A. & Wasserman, L. (1997). Computing Bayes factors by combining simulation and asymptotic approximations. Journal of the American Statistical Association 92 903–915.

[27] Ding, D., & and R. C. Sickles (2018). Frontier efficiency, capital structure, and portfolio risk: An empirical analysis of U.S. banks. BRQ Business Research Quarterly 21 (4), 262–277.

[28] Du, K., A. C. Worthington, & V. Zelenyuk, 2018. Data envelopment analysis, truncated regression and double-bootstrap for panel data with application to Chinese banking. European Journal of Operational Research 265 (2), 748–764.

[29] Durham, G., & J. Geweke, 2013, Adaptive Sequential Posterior Simulators for Massively Parallel Computing Environments. Advances in Econometrics (Ivan Jeliazkov and Dale J. Poirier, eds.). Bayesian Model Comparison (Advances in Econometrics, Volume 34) Emerald Group Publishing Limited. Chapter 1, 1–44.

[30] Fogliatto, F.S., & Anzanello, M.J., 2011. Learning curves: the state of the art and research directions. In: Jaber, M.Y. (Ed.), Learning Curves: Theory, Models, and Applications. CRC, Boca Raton, pp. 3–21.

[31] Fong, E., & Holmes, 2020, On the marginal likelihood and cross-validation, Biometrika, 107 (2), 489–496.

[32] Fukuyama, F., & R. Matousek, 2017. Modelling bank performance: A network DEA approach. European Journal of Operational Research 259 (2), 721–732.

[33] Gadzinski, G., & Castello, A., 2020, Fast and Frugal heuristics augmented: When machine learning quantifies Bayesian uncertainty. Journal of Behavioral and Experimental Finance 26, 100293, in press.

[34] Ge, L., & Peng, X.M., 2012. Research fronts of international entrepreneurship education in the visual threshold of knowmetrics. J. Knowl. Bas. Innov. China 4 (1), 55–65.

[35] Geweke, J., 1992. Evaluating the accuracy of sampling-based approaches to calculating posterior moments. In *Bayesian Statistics 4* (J. M. Bernado, J. O. Berger, A. P. Dawid and A. F. M. Smith, eds.). Clarendon Press, Oxford, UK, 169–193.

[36] Gilboa, I., Postlewaite A., & Schmeidler, D. (2008). Probability and uncertainty in economic modelling, J. Econ. Perspectives 22, 173–188.

[37] Gilboa, I., Postlewaite, A., & Schmeidler, D. (2009). Is it always rational to satisfy Savage's axioms?, Econ. and Phil. 25, 285–296.

[38] Gilks, W. R., & C. Berzuini, 2001. Following a moving target: Monte Carlo inference for dynamic Bayesian models, Journal of the Royal Statistical Society B 63, 127–146.

[39] Giustozzi, F., J. Sauniera, & C. Zanni-Merk (2018). Context Modeling for Industry 4.0: an Ontology-Based Proposal. Procedia Computer Science 126, 675–684.

[40] Glock, C.H., Grosse, E.H., Jaber, M.Y., & Smunt, T.L., 2019. Applications of learning curves in production and operations management: a systematic literature review. Comput. Ind. Eng. 132, 422–441.

[41] Golub, B., and Jackson M. (2010). Naïve Learning in Social Networks and the Wisdom of Crowds, Amer. Econ. J. Microeconomics 2 (1), 112–149.

[42] Greene, W. H., 1980a. Maximum likelihood estimation of econometric frontier functions. J. Econometrics 13 (1), 27–56.

[43] Greene, W. H., 1980b. On the estimation of a flexible frontier production model. J. Econometrics 13 (1), 101–115.

[44] Greene, W. H., 1990. A Gamma-distributed stochastic frontier model. J. Econometrics 46 (1–2), 141–163.

[45] Grosse, E.H., & Glock, C.H., 2015. The effect of worker learning on manual order picking processes. Int. J. Prod. Econ. 170, 882–890.

[46] Grosse, E.H., Glock, C.H., & Muller, S., 2015. Production economics and the learning curve: a meta-analysis. Int. J. Prod. Econ. 170, 401–412.

[47] Hautsch, N., & Hess, D. (2007). Bayesian Learning in Financial Markets: Testing for the Relevance of Information Precision in Price Discovery. The Journal of Financial and Quantitative Analysis 42 (1), 189–208.

[48] Hornik, K., M. Stinchcombe, & H. White (1989). Multi-layer feedforward networks are universal approximators, Neural Networks 2, 359–366.

[49] Jaber, M.Y., 2006. Learning and forgetting models and their applications. In: Badiru, A.B. (Ed.), Handbook of Industrial And Systems Engineering. CRC Press, Boca Raton, ISBN 9781466515048, pp. 30.1–30.24 (Reprinted in 2013 in Handbook of Industrial and Systems Engineering, 2nd edition, pages 535–566, Boca Raton: CRC Press.

[50] Jaber, M.Y., Bonney, M., 1999. The economic manufacture/order quantity (EMQ/EOQ) and the learning curve: past, present, and future. Int. J. Prod. Econ. 59 (1), 93–102.

[51] Jaber, M.Y., & Guiffrida, A.L., 2004. Learning curves for processes generating defects requiring reworks. Eur. J. Oper. Res. 159 (3), 663–672.

[52] Jaber, M.Y., & Khan, M., 2010. Managing yield by lot splitting in a serial production line with learning, rework and scrap. Int. J. Prod. Econ. 124 (1), 32–39.

[53] Jaber, M. Y., Peltokorpi, J., Glock, C. H., Grosse, & E. H., Pusic, M. (2021). Adjustment for cognitive interference enhances the predictability of the power learning curve. International Journal of Production Economics 234, 108045, in press.

[54] Johnson-Laird, P. N. (1983). Mental Models: Towards a Cognitive Science of Language, Inference, and Consciousness. Cambridge, MA: Harvard University Press.

[55] Jones, M., & Love, B. C. (2011). Bayesian Fundamentalism or Enlightenment? On the explanatory status and theoretical contributions of Bayesian models of cognition. Behavioral and Brain Sciences 34 (4), 169–188.

[56] Jordan, J. S., 1991. Bayesian learning in normal form games. Games and Economic Behavior, 3 (1), 60–81.

[57] Jovanovic, B., & Nyarko, Y., 1996. Learning by doing and the choice of technology. Econometrica 64 (6), 1299–1310.

[58] Kahneman, D. (2011) Thinking, Fast and Slow. London: Penguin.

[59] Kahneman D., & Tversky, A. (1974). Judgement under uncertainty: heuristics and biases, Science 185, 1124–1131.

[60] Kass, R.E. & Raftery, A.E., (1995), Bayes factors, Journal of the American Statistical Association, 90 (430), 773–795.

[61] Kemp, C., Perfors, A. & Tenenbaum, J. B. (2007) Learning overhypotheses with hierarchical Bayesian models. Developmental Science (10), 307–321.

[62] Kevork, I. S., J. Pange, P. Tzeremes, & N. G. Tzeremes, 2017. Estimating Malmquist productivity indexes using probabilistic directional distances: An application to the European banking sector. European Journal of Operational Research 261 (3), 1125–1140.

[63] Koetter, M., J. W. Kolari, & L. Spierdijk, L., 2012. Enjoying the quiet life under deregulation? Evidence from adjusted Lerner indices for U.S. banks, The Review of Economics and Statistics 94 (2), 462–480.

[64] Koetter, M., & O.Vins (2008). The quiet life hypothesis in banking-evidence from German Savings banks. Department of Finance, Goethe University, Working Paper Series: Finance and Accounting No. 190, Frankfurt.

[65] Lampe, H. W., & D. Hilgers, 2015. Trajectories of efficiency measurement: A bibliometric analysis of DEA and SFA. European Journal of Operational Research 240 (1), 1–21.

[66] Lucas, R. E., Jr. (1973). Some International Evidence on Output-Inflation Tradeoffs. American Economic Review 63, 326–334,

[67] Malikov, E., S. C. Kumbhakar, & M. G. Tsionas, 2016. A cost system approach to the stochastic directional technology distance function with undesirable outputs: the case of U.S. banks in 2001-2010. Journal of Applied Econometrics, 31 (7), 1407–1429.

[68] Mercier, H., & Sperber, D. (2017). The Enigma of Reason: A New Theory of Human Understanding. London: Penguin Random House.

[69] Minniti, M., & Bygrave, W., 2001, A Dynamic Model of Entrepreneurial Learning, Entrepreneurship Theory and Practice 25 (3), 5–16.

[70] Molavi, P., Tahbaz-Salehi, A., & Jadbabaie, A. (2018). A Theory of Non-Bayesian Social Learning. Econometrica 86 (2), 445–490.

[71] Moreno, A.M., & Casillas, J.C., 2008. Entrepreneurial orientation and growth of SMEs: a causal model. Entrep. Theory Pract. 32 (3), 507–528.

[72] Nachbar, J. H. 1997. Prediction, Optimization, and Learning in Repeated Games. Econometrica 65, 275–309.

[73] Nachbar, J. H., . 2005. Beliefs in Repeated Games. Econometrica 73, 459–480.

[74] Oaksford, M., & Chater, N. (2007). Bayesian rationality: The probabilistic approach to human reasoning. Oxford: Oxford University Press.

[75] Oikawa, K., 2016. A microfoundation for stochastic frontier analysis. Economics Letters 139, 15–17.

[76] Paul, S., & Shankar, S., 2018. On estimating efficiency effects in a stochastic frontier model European Journal of Operational Research, 271, 769–774.

[77] Plaza, M., & Rohlf, K., 2008. Learning and performance in ERP implementation projects: a learning-curve model for analyzing and managing consulting costs. Int. J. Prod. Econ. 115 (1), 72–85.

[78] Pusic, M.V., Boutis, K., Hatala, R., & Cook, D.A., 2015. Learning curves in health professions education. Acad. Med. 90 (8), 1034–1042. https://doi.org/10.1097/ ACM.0000000000000681.

[79] Pusic, M.V., Boutis, K., Pecaric, M.R., Savenkov, O., Beckstead, J.W., & Jaber, M.Y., 2017. A primer on the statistical modelling of learning curves in health professions education. Adv. Health Sci. Educ. 22 (3), 741–759.

[80] Rabin M., 1998. Psychology and economics, J. Econ. Lit. 36, 11–46.

[81] Rauch, A., Wiklund, J., Lumpkin, G.T., & Frese, M., 2009. Entrepreneurial orientation and business performance: an assessment of past research and suggestions for the future. Entrep. Theory Pract. 33 (3), 761–787.

[82] Roy, V., Silvestre, B. S., & Singh, S., 2020. Reactive and proactive pathways to sustainable apparel supply chains: Manufacturer's perspective on stakeholder salience and organizational learning toward responsible management. International Journal of Production Economics 227, 107672. in press.

[83] Sabahi, S., & Parast, M. M., 2020. The impact of entrepreneurship orientation on project performance: A machine learning approach. International Journal of Production Economics 226, 107621, in press.

[84] Schreiber, G., Akkermans, H., Anjewierden, A., Hoog, R.D., Shadbolt, N.R., & Wielinga, B. (2000). Knowledge Engineering and Management: The CommonKADS Methodology; vol. 99. ISBN 0262193000. doi:10.1016/S0933-3657(01)00090-2. URL http://eprints.ecs.soton.ac.uk/2270/

[85] Stamboulis, Y., & Barlas, A., 2014. Entrepreneurship education impact on student attitudes. Int. J. Manag. Educ. 12 (3), 365–373.

[86] Tortorella, G. L., Vergara, A. M. C., Garza-Reyes, J. A., & Sawhney, R., 2020, Organizational learning paths based upon industry 4.0 adoption: An empirical study with Brazilian manufacturers. International Journal of Production Economics 219, 284–294.

[87] Tsionas, M. G., 2000. Full likelihood inference in normal-gamma stochastic frontier models. Journal of Productivity Analysis 13, 183–205.

[88] Tsionas, M. G., 2017, Microfoundations for stochastic frontiers. European Journal of Operational Research 258, 1165–1170.

[89] Tsionas, M. G., 2020, Bounded rationality and thick frontiers in stochastic frontier analysis. European Journal of Operational Research 284, 762–768.

[90] Tsionas, M. G., E. Malikov, & S. C. Kumbhakar, 2018. An internally consistent approach to the estimation of market power and cost efficiency with an application to U.S. banking. European Journal of Operational Research Volume 270 (2), 747–760.

[91] Tsionas, M. G., & E. C. Mamatzakis, 2017. Adjustment costs in the technical efficiency: An application to global banking. European Journal of Operational Research 256 (2), 640–649.

[92] Tsionas, M. G., Mamatzakis, E. (2019). Further results on estimating inefficiency effects in stochastic frontier models. European Journal of Operational Research 275 (3), 1157-1164.

[93] Tsionas, M. G., E. Mamatzakis, & S. Ongena, 2020. Does risk aversion affect bank output loss? The case of the Eurozone. European Journal of Operational Research 282 (3), 1127–1145.

[94] Ubøe, J., J. Andersson, K. Jörnsten, J. Lillestøl, & L. Sandal, 2017, Statistical testing of bounded rationality with applications to the newsvendor model. European Journal of Operational Research 259 (1), 251–261.

[95] Wang, J., & Xiao, B., 2017. A minmax regret price control model for managing perishable products with uncertain parameters. European Journal of Operational Research, 258 (2), 652–663.

[96] White, H., 1989. Learning in artificial neural networks: A statistical perspective. Neural Computation 1, 425–464.

[97] White, H., 1990. Connectionist nonparametric regression: Multilayer feedforward networks can learn arbitrary mappings, Neural Networks 3, 535–549.

[98] Xiao, B. & Yang, W., 2021. A Bayesian learning model for estimating unknown demand parameter in revenue management. European Journal of Operational Research 293 (1), 248–262.