



HAL
open science

Revisiting Identification Concepts in Bayesian Analysis

Jean-Pierre Florens, Anna Simoni

► **To cite this version:**

Jean-Pierre Florens, Anna Simoni. Revisiting Identification Concepts in Bayesian Analysis. *Annals of Economics and Statistics*, 2021, 144, pp.1-38. 10.15609/annaeconstat2009.144.0001 . hal-03504692

HAL Id: hal-03504692

<https://hal.science/hal-03504692>

Submitted on 29 Dec 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Revisiting identification concepts in Bayesian analysis *

Jean-Pierre Florens[†] Anna Simoni[‡]

October 19, 2021

Abstract

This paper studies the role played by identification in the Bayesian analysis of statistical and econometric models. First, for unidentified models we demonstrate that there are situations where the introduction of a non-degenerate prior distribution can make a parameter that is nonidentified in frequentist theory identified in Bayesian theory. In other situations, it is preferable to work with the unidentified model and construct a Markov Chain Monte Carlo (MCMC) algorithms for it instead of introducing identifying assumptions. Second, for partially identified models we demonstrate how to construct the prior and posterior distributions for the identified set parameter and how to conduct Bayesian analysis. Finally, for models that contain some parameters that are identified and others that are not we show that marginalizing out the identified parameter from the likelihood with respect to its conditional prior, given the nonidentified parameter, allows the data to be informative about the nonidentified and partially identified parameter. The paper provides examples and simulations that illustrate how to implement our techniques.

Key words: Minimal Sufficiency, Exact Estimability, Set identification, Dirichlet Process, Capacity functional, Nonparametric models.

JEL code: C11; C10; C14

*First draft: February 2011. The authors gratefully thank the Editor Laurent Linnemer and two anonymous referees for their constructive comments on the previous version of the paper. We acknowledge helpful comments from Andrew Chesher, Frank Kleibergen and participants to the 2011-Cowles Foundation annual Econometrics Conference, ESEM-2011 in Oslo and Econometric seminar in GREQAM-Marseille. Anna Simoni gratefully acknowledges financial support from Labex ECODEC (ANR-11-LABEX-0047).

[†]Toulouse School of Economics, Université de Toulouse Capitole, Toulouse - 21, allée de Brienne - 31000 Toulouse (France). Email: jean-pierre.florens@tse-fr.eu

[‡]CREST, CNRS, ENSAE, École Polytechnique - 5, avenue Henry Le Chatelier, 91120 Palaiseau, France. Email: anna.simoni@ensae.fr (*corresponding author*).

1 Introduction

This paper investigates Bayesian analysis in models that lack identification. We first revisit theoretical concepts related to identification. We highlight how lack of identification has a different impact for statistical inference depending on whether one develops Bayesian or frequentist inference. Specifically, Bayesian analysis can be carried out without imposing any identifying restrictions. Then, we distinguish between nonidentified and partially identified models. For partially identified models we propose to construct the prior and posterior distributions of a set parameter by using the capacity functional and we introduce the concepts of *prior* (resp. *posterior*) *capacity functional* and *prior* (resp. *posterior*) *coverage function*.

As Lindley (1971) remarked, the problem of non identification causes no real difficulty in the Bayesian approach. Indeed, if a proper prior distribution is specified, then the posterior distribution is well-defined. Kadane (1974) observed that identification is a property of the likelihood function which is the same irrespective of whether it is considered from Bayesian or frequentist perspectives. It is however necessary to distinguish between three concepts of identification depending on the level of specification of the model: *sampling* (or *frequentist*) *identification*, *measurable identification* and *Bayesian identification*, see *e.g.* Florens et al. (1985, 1990) and Florens and Mouchart (1986). Sampling identification is defined without the introduction of a σ -field associated with the parameter space. On the other hand, such a σ -field is necessary for the other two concepts of identification. In addition, the notion of Bayesian identification requires the introduction of a unique joint probability measure over the sample and the parameters.

When a σ -field associated with the parameter space is introduced, the concept of identification is related to the *minimal sufficient* parameter. Namely, the observed sample brings information only on the *minimal sufficient* parameter and hence, the parameter of the model is identified if it is equal to the minimal sufficient parameter. In other words, the *minimal sufficient* parameter is the smallest σ -field on the parameter space that makes the sampling probability measurable. So, it is the identified parameter. Conditionally on this parameter, the Bayesian experiment is completely non informative: the prior distribution of a nonidentified parameter is not revised through the information brought by the data so that the conditional posterior and conditional prior distributions (conditioned on the identified parameter)

are the same.

Equivalently, we say that a model is nonidentified if the parametrization is redundant. It is then natural to wonder why one should introduce a σ -field larger than the minimal sufficient parameter. In nonexperimental fields, redundant parametrization is usually introduced either as an early stage of model building or as a support for relevant prior information or because the parameter of interest (making *e.g.* the loss function measurable) is larger than the minimal sufficient parameter (see **Example 1**). In experimental fields, it may be the case that the experimental design will not provide information on all the parameters of a theoretically relevant model, see Florens et al. (1990).

This paper makes the following contributions. (I) For nonidentified models, we show that there are situations where the introduction of a non-degenerate prior distribution can make a parameter that is unidentified in frequentist theory identified in Bayesian theory. Specifically, we demonstrate that this is true for nonparametric models with heterogeneity modeled either as a Gaussian process or as a Dirichlet process where the parameter of interest is the (hyper)parameter of the heterogeneity distribution. We show that in these models it is possible to obtain Bayesian identification since the hyperparameter can be expressed as a known function of the identified parameter. We stress that this is not a property of the prior of the nonidentified parameter, but instead it is a property of the conditional prior of the identified parameter, given the unidentified parameter. Such a result is no longer true in parametric models where a degenerate prior for the nonidentified parameter is required to get Bayesian identification, which then is completely artificial. (II) We provide the example of latent variable models to illustrate that it is preferable to conduct Bayesian inference and develop Markov Chain Monte Carlo (MCMC) algorithms for the nonidentified model instead of introducing identifying restrictions. Working with the nonidentified model grants better mixing properties of the MCMC.

(III) We propose a procedure to make Bayesian analysis for partially identified models where the identified parameter is a set, called the *identified set*. For these models we build up a new Bayesian nonparametric approach, based on the Dirichlet process prior, and construct prior and posterior distributions for set parameters. We propose to define these distributions in terms of *prior* and *posterior capacity functionals*. The posterior capacity functional is an appealing tool to build estimators and credible sets for the identified set and, in addition, it can be easily computed either by simulations or in closed-form. (IV) We show that, when the model contains an identified parameter and a parameter

of interest that lacks identification, the latter can be identified in the marginal model where the identified parameter has been marginalized out from the likelihood function with respect to its conditional prior given the unidentified parameter. Therefore, the integrated likelihood depends on the unidentified parameter and the prior of the latter is revised by the information brought by the data in the marginal model. To complement our theoretical approach we develop many examples and simulations that show how our method can be implemented.

Our results contribute to show that the Bayesian approach is appealing in models that lack identification for several reasons. First, if the prior distribution on all the parameters is proper, Bayesian analysis of nonidentified and partially identified models is always possible since the posterior distribution always exists. Second, when the parameter in the model is multidimensional, with some components that are identified and others that are nonidentified, data can be marginally informative about the nonidentified parameter. That is, if the parameters are a priori dependent, then after we marginalise out the identified parameter from the likelihood with respect to its conditional prior given the nonidentified parameter, the integrated likelihood will depend on the nonidentified parameter. Third, the issue of non- and/or partial-identification can be reduced (or even eliminated) by introducing an informative prior. Lastly, even when the model is nonidentified or partially identified, Bayesian procedures have computational advantages over frequentist ones. In particular, MCMC algorithms have better mixing properties if they are specified for the nonidentified model instead of imposing identifying restrictions.

The paper is organized as follows. In section 2 we discuss the three concepts of identification given above and the concept of partial identification. Section 3 studies models with heterogeneity and models with latent variables. Bayesian analysis of partially identified models is developed in section 4. Finally, in section 5 we discuss identification by marginalisation for both unidentified and partially identified models. Section 6 concludes. Examples and all the proofs are in Appendix E in the Supplement. In the paper we abbreviate “almost surely” by “a.s.”, P will denote the data distribution and p the associated Lebesgue density. For an event A , $\mathbb{1}\{A\}$ denotes the indicator function which takes the value 1 if the event A is satisfied and 0 otherwise.

Literature Review. Our paper is related to two strands of the Bayesian literature, the one focusing on nonidentified models and the Bayesian literature on partially identified models. We provide here a

concise review of the previous contributions that are the most relevant for our paper.

Initial discussions of nonidentified models which lay the foundations of nonidentification in a Bayesian experiment in a measure-theoretic framework can be found in Lindley (1971), Kadane (1974) and Picci (1977), among others. Florens et al. (1985, 1990) and Florens and Mouchart (1986) resume these works and provide further developments. In particular, they provide a rigorous discussion on the difference among *sampling*, *measurable* and *Bayesian identification*. Compared to these contributions, in section 2 we provide a unified framework that gather together the concepts and results related to nonidentification that are the most relevant for applied Bayesian analysis in econometrics. We adopt a measure-theoretic framework and present the results of identification in terms of σ -fields of sets of the parameter space. An aspect that we do not consider in our paper is the prior elicitation process for nonidentified parameters which is investigated for instance in San Martín and González (2010).

As in Poirier (1998), we emphasise the difference between marginal and conditional uninformativeness of the data for Bayesian analysis of nonidentified models. The main contribution of Poirier (1998) consists in analysing the diverse effect of nonidentification for Bayesian analysis in the two following situations: the case with proper priors, and the case with improper priors. Our paper does not emphasize this difference between proper and improper priors. Instead, one of our main contributions is to demonstrate the different role played by the prior in nonparametric and parametric models that are nonidentified. The question of nonidentification in nonparametric models has not been explicitly considered in the past Bayesian literature to the best of our knowledge. For these models we prove that Bayesian identification, *i.e.* through the prior, arises in a non-artificial way in some cases.

Gustafson (2005) analyses nonidentifiability by taking a nonconventional approach. Instead of contracting the model, which consists in reducing the redundant parametrisation to get identification, he proposes to expand the model to a supermodel that can at best yield identification by the prior. This is an alternative approach to ours.

The second strand of literature that relates to our paper is the Bayesian literature on partial identification. It includes relatively few contributions in comparison to the vast frequentist literature on partially identified models. For excellent overviews of frequentist and Bayesian inference for partially identified econometric models see Bontemps and Magnac (2017), Molinari (2020) and references therein. Most of the previous Bayesian literature is interested in constructing Bayesian procedures that

provide asymptotically valid frequentist inferences for partially identified models, see for instance Liao and Jiang (2010), Gustafson (2010, 2012), Moon and Schorfheide (2012), Norets and Tang (2013), Kline and Tamer (2016), Chen et al. (2018), Liao and Simoni (2019) and Giacomini and Kitagawa (2021). They have proposed Bayesian or quasi-Bayesian approaches for constructing (asymptotically) valid frequentist confidence sets of the identified set. Their methods are mostly based either on the limited information likelihood of Chernozhukov and Hong (2003) or on the support function. Compared to this literature, we take a different approach that is based on Dirichlet process prior placed on the identified reduced-form parameter. We construct the prior and posterior probabilities for the identified set as prior and posterior capacity functionals, which is new in the literature. In addition, we are not interested in frequentist asymptotic properties of our procedure and conduct a fully Bayesian analysis.

Marginal analysis in partially identified models, which we consider in section 5, has been considered in Liao and Jiang (2010) but in a way different from ours. They marginalise out the slackness parameters in the partially identified models characterized by moment inequalities while we marginalise out the identified (reduced-form) parameter.

2 Some General Definition

2.1 Sampling Identification

In this section we recall basic definitions in the non-Bayesian framework, called in the paper sampling theory approach. Let us consider a statistical model defined by a sampling space X provided with a σ -field \mathcal{X} and by a collection of probabilities on this space. The collection of sampling probability measures defined on (X, \mathcal{X}) is denoted by $(P^\theta)_{\theta \in \Theta}$ and is in general indexed by a parameter $\theta \in \Theta$, which can be finite dimensional or a functional parameter. Hence, the *sampling statistical model* \mathcal{E}_s is defined as $\mathcal{E}_s := \{\Theta, (X, \mathcal{X}), (P^\theta)_{\theta \in \Theta}\}$.

In a small sample approach we observe a finite number of realizations of a random variable taking values in a measurable space – for example an *iid* sample $x := (x_1, \dots, x_n) \in (X, \mathcal{X})$ – where the sample size is not made explicit and is kept fixed. In an asymptotic approach, \mathcal{X} is provided with a filtration $(\mathcal{X}_n)_{n \geq 1}$, where \mathcal{X}_n represents the information contained in a sample of size n . We denote by

$\mathcal{X}_\infty := \bigvee_{n \geq 1} \mathcal{X}_n$ the σ -field generated by $\bigcup_{n \geq 1} \mathcal{X}_n$ and write $\mathcal{X}_n \rightarrow \mathcal{X}_\infty$ as $n \rightarrow \infty$.

Two parameters θ_1 and θ_2 are said to be *observationally equivalent*, in symbol $\theta_1 \sim \theta_2$, if $P^{\theta_1} = P^{\theta_2}$. This relation defines an equivalence relation and we denote by $\tilde{\Theta}$ the quotient space Θ / \sim , i.e. the elements $\tilde{\theta}$ of $\tilde{\Theta}$ are the equivalence classes over Θ by \sim . Notice that $\tilde{\theta}$ is a set and the elements of $\tilde{\Theta}$ are sets of parameters. We now define the concept of sampling identification.

Definition 2.1 (Sampling identification) *A real valued function $a(\theta)$ defined on Θ is identified if $\theta_1 \sim \theta_2$ implies $a(\theta_1) = a(\theta_2)$. The sampling model \mathcal{E}_s is identified if and only if any real valued function defined on Θ is identified.*

This definition is equivalent to say that $\theta_1 \sim \theta_2$ implies $\theta_1 = \theta_2$ or to say that any equivalence class is reduced to a singleton. In turn, this is equivalent to say that the sampling model \mathcal{E}_s is identified if the mapping $\theta \rightarrow P^\theta$ is injective.

For any sampling statistical model \mathcal{E}_s there exists a canonical identified model $\tilde{\mathcal{E}}_s$ defined as the sampling statistical model with parameter space $\tilde{\Theta} : \tilde{\mathcal{E}}_s = \{\tilde{\Theta}, (X, \mathcal{X}), (P^{\tilde{\theta}})_{\tilde{\theta} \in \tilde{\Theta}}\}$ where $P^{\tilde{\theta}} = P^\theta$ for any $\theta \sim \tilde{\theta}$. Equivalently, $\tilde{\mathcal{E}}_s$ is the set identified statistical model associated with \mathcal{E}_s . Thus, one may construct an identified sampling model by selecting a single element in each equivalence class. Let us call *section* a function $\sigma : \tilde{\Theta} \rightarrow \Theta$ such that $\sigma(\tilde{\theta}) \in \tilde{\theta}$ and $\sigma(\tilde{\theta}) \in \Theta$. By using such a function one may define the identified sampling model

$$\mathcal{E}_{s,\sigma} := \{\Theta_\sigma, (X, \mathcal{X}), (P^\theta)_{\theta \in \Theta_\sigma}\}, \quad \Theta_\sigma := \sigma(\tilde{\Theta}),$$

where Θ_σ is the image of $\tilde{\Theta}$ by σ and $P^\theta = P^{\sigma(\tilde{\theta})}$. In this context it is natural to look for continuous sections $\sigma(\cdot)$ or to bicontinuous bijections between $\tilde{\Theta}$ and Θ_σ . In general, it is not possible to devise constructive rules for selecting elements in Θ for each $\tilde{\theta} \in \tilde{\Theta}$. The existence of Θ_σ cannot be proved by using axioms from set theory if we do not have some structure on the parameter space but it must be asserted as an additional axiom called *axiom of choice*, see e.g. Kolmogorov and Fomin (1975).

In the sampling theory statistics, a topological structure is needed in the parameter space in order to define statistical decision rules based on convergence, risk or loss function. A canonical topological structure is defined on Θ and then may be carried on $\tilde{\Theta}$. Let $\rho : \Theta \rightarrow \tilde{\Theta}$ be the canonical application $\theta \rightarrow \rho(\theta) = \tilde{\theta}$. The natural topological structure is the smallest one for which ρ is continuous, i.e.

$\rho^{-1}(\tilde{O})$ is open in Θ whenever $\tilde{O} \subset \tilde{\Theta}$ is open. We will not detail the topological aspects of set identification, which is beyond the scope of this paper, and refer to Husmoller (1994) and appendix to chapter 3 in Dellacherie and Meyer (1975).

2.2 Measurable Identification

Let us define a measurable statistical model \mathcal{E}_m as $\mathcal{E}_m := \{(\Theta, \mathcal{A}), (X, \mathcal{X}), (P^\theta)_{\theta \in \Theta}\}$, where we use the notation previously introduced and in addition we define \mathcal{A} to be a σ -field on Θ such that P^θ is a transition probability. Recall that given two measurable spaces (Θ, \mathcal{A}) and (X, \mathcal{X}) , the mapping $P^{(\cdot)}(\cdot) : \Theta \times \mathcal{X} \rightarrow [0, 1]$ is a transition probability if: (i) $\forall \theta \in \Theta$, $P^\theta(\cdot)$ is a probability measure on (X, \mathcal{X}) , and (ii) $\forall E \in \mathcal{X}$, $P^{(\cdot)}(E)$ is a measurable function on (Θ, \mathcal{A}) . The introduction of the σ -field \mathcal{A} of subsets of Θ is necessary in order to introduce a joint probability measure on the product space $\Theta \times X$. This probability will be introduced in section 2.3. In this sense, the introduction of a measurable statistical model is a preliminary step for the construction of a Bayesian model.

A σ -field represents an information structure and the parameter of interest is naturally introduced as a σ -field \mathcal{A} that makes both the transition probability $P^{(\cdot)}(E)$, $\forall E \in \mathcal{X}$, and the loss function of the underlying decision model, \mathcal{A} -measurable. We now introduce the concept of sufficient σ -field which is essential in order to discuss identification.

Definition 2.2 (Sufficiency and minimal sufficiency) *A sub- σ -field \mathcal{B} of \mathcal{A} is said to be sufficient in the model \mathcal{E}_m if $P^{(\cdot)}(E)$ is \mathcal{B} -measurable for any $E \in \mathcal{X}$. A sub- σ -field \mathcal{B} of \mathcal{A} is said to be minimal sufficient if: (i) \mathcal{B} is sufficient and (ii) $\mathcal{C} \subset \mathcal{A}$ sufficient implies that $\mathcal{B} \subset \mathcal{C}$.*

The structure of model \mathcal{E}_m allows for the possibility of introducing a given distribution on Θ . Here, we discuss identification in model \mathcal{E}_m without the specification of a (prior) distribution on (Θ, \mathcal{A}) which will be discussed in the next section. The following proposition introduces the concept of identification in model \mathcal{E}_m , called *measurable identification*.

Proposition 2.1 (Measurable identification) *There exists a minimal sufficient σ -field \mathcal{A}_* equal to the intersection of all the sufficient σ -fields. The measurable model \mathcal{E}_m is identified if and only if $\mathcal{A}_* = \mathcal{A}$.*

Definition 2.3 (Measurable identification of a function) *A real valued function $a(\cdot)$ defined on (Θ, \mathcal{A}) is measurably identified if it is \mathcal{A}_* -measurable.*

The notion of sampling and measurable identification are identical if Θ is a measurable subset of \mathbb{R}^k provided with the Borelian σ -field and if X is included in \mathbb{R}^n and also provided with the Borelian σ -field. This equivalence is true more generally and requires that Θ “looks like” a Borelian of \mathbb{R} and that the σ -field on X is separable or equivalently generated by a countable family of subsets. Such a link between sampling and measurable identification is established in the next theorem where we assume that (Θ, \mathcal{A}) is a Souslin space. We recall that a measurable space (Θ, \mathcal{A}) is a Souslin space if there exists an analytic set $B \subset \mathbb{R}$ and a bimeasurable bijection between (Θ, \mathcal{A}) and $(B, B \cap \mathcal{B})$, where $B \cap \mathcal{B}$ denotes the restriction to B of the Borel σ -fields \mathcal{B} of \mathbb{R} . An analytic set on \mathbb{R} is the projection on \mathbb{R} of a Borelian set in \mathbb{R}^2 . In particular, all Borelian sets are analytic. The property that (Θ, \mathcal{A}) is a Souslin space is clearly true for finite dimensional parameter spaces or more generally for Polish spaces, which include the L^2 spaces on a real space. The majority of the functional spaces usually considered in statistical and econometric applications are Polish, so the requirement that (Θ, \mathcal{A}) is a Souslin space is almost always satisfied.

Theorem 2.1 *Let us assume that (Θ, \mathcal{A}) is a Souslin space. Then,*

- (1) *a real valued function $a(\cdot)$ defined on (Θ, \mathcal{A}) is measurable identified if and only if $a(\cdot)$ is constant on the equivalence class $\tilde{\Theta}$;*
- (2) *if \mathcal{X} is separable, a model is sampling identified if and only if it is measurable identified.*

The first point of the theorem provides an additional characterisation of measurable identification. The second part establishes that sampling and measurable identification are equivalent when \mathcal{X} is separable.

2.3 Bayesian Identification

A Bayesian model consists of a measurable statistical model and a measure on (Θ, \mathcal{A}) which can be either proper (if it is a probability measure) or improper, called *prior distribution* and denoted by μ . For simplicity we will consider a probability measure. Then, μ and P^θ generate a unique measure on $(\Theta \times X, \mathcal{A} \otimes \mathcal{X})$ denoted by Π :

$$\Pi(A \times E) = \int_A P^\theta(E) \mu(d\theta) = \int_E \mu^x(A) P(dx), \quad A \in \mathcal{A}, E \in \mathcal{X}$$

or, equivalently $\Pi = \mu \otimes P^\theta = P \otimes \mu^x$, where P is the marginal probability measure on (X, \mathcal{X}) called the predictive distribution and μ^x is the posterior distribution. We assume that there exists a regular version of the conditional probability on Θ given X , that is, $\mu^x(\cdot)$ is a transition probability (see definition in section 2.2). The Bayesian model is then defined by the following probability space

$$\mathcal{E}_b := \{\Theta \times X, \mathcal{A} \otimes \mathcal{X}, \Pi\}.$$

In the following, we may use the notation $\mu(\cdot|x)$ instead of $\mu^x(\cdot)$ when it is more appropriate. Moreover, by abuse of notation we use $\mu(\theta|x)$ (resp. $\mu(\theta)$) to denote both the posterior (resp. the prior) distribution and its Lebesgue density function. The sampling (resp. predictive) density function with respect to the Lebesgue measure is denoted by $p(x|\theta)$ (resp. $p(x)$). Before defining the concept of Bayesian identification, we recall the definition of sufficiency and minimal sufficiency in the Bayesian model.

Definition 2.4 (Bayesian sufficiency) *A sub σ -field \mathcal{B} of \mathcal{A} is sufficient in the Bayesian model if and only if $\mathcal{A} \perp \mathcal{X}|\mathcal{B}$.*

The conditional independence of the previous definition has two equivalent characterizations: (i) for every positive function $t : \mathcal{X} \rightarrow \mathbb{R}_+$, $\mathbf{E}(t(x)|\mathcal{A}) = \mathbf{E}(t(x)|\mathcal{B})$, $\Pi - a.s.$ provided that the conditional expectations exist; (ii) for every positive function $a : \mathcal{A} \rightarrow \mathbb{R}_+$, $\mathbf{E}(a(\theta)|\mathcal{X} \vee \mathcal{B}) = \mathbf{E}(a(\theta)|\mathcal{B})$, $\Pi - a.s.$

The first characterization (i) weakens the concept of sufficient σ -field because the property is only required almost surely with respect to the prior probability. If \mathcal{A} is generated by a function $a(\cdot)$ and \mathcal{B} by a function $b(\cdot)$ then (i) means that the likelihood functions $p(x|a)$ and $p(x|b)$ are *a.s.* equal. The second characterization (ii) says that the conditional prior and posterior distributions are *a.s.* equal given a sufficient σ -field. Equivalently, (ii) means that the posterior distribution $\mu(a|x, b)$ of a is *a.s.* equal to the prior $\mu(a|b)$.

It may be proven that there exists a *minimal Bayesian sufficient σ -field* in \mathcal{A} , denoted by \mathcal{A}_*^μ and defined as follows.

Definition 2.5 (Minimal Bayesian sufficiency) *A minimal Bayesian sufficient σ -field \mathcal{A}_*^μ (also called μ -a.s. minimal sufficient σ -field) is the σ -field generated by all the versions of $\mathbf{E}(t(x)|\mathcal{A})$ for any positive \mathcal{X} -measurable function $t : \mathcal{X} \rightarrow \mathbb{R}_+$, provided the expectation exists. This σ -field \mathcal{A}_*^μ is called*

the projection of \mathcal{X} on \mathcal{A} and is also denoted by $\mathcal{A}\mathcal{X}$.

The minimal Bayesian sufficient σ -field \mathcal{A}_*^μ is a σ -field on the parameter space and depends on the prior μ . An equivalent definition of \mathcal{A}_*^μ is that \mathcal{A}_*^μ is the smallest σ -field which makes the sampling probabilities $P^{(\cdot)}(E)$ measurable, $\forall E \in \mathcal{X}$, completed by the null sets of \mathcal{A} with respect to μ . It may be easily verified that $\overline{\mathcal{A}_*} \cap \mathcal{A} = \mathcal{A}\mathcal{X} \equiv \mathcal{A}_*^\mu$ where $\overline{\mathcal{A}_*}$ is the σ -field generated by \mathcal{A}_* and all the null sets of $\mathcal{A} \vee \mathcal{X}$, and $\overline{\mathcal{A}_*} \cap \mathcal{A}$ is the σ -field generated by \mathcal{A}_* and all the null sets of \mathcal{A} with respect to μ .

We are now ready to give the definition of Bayesian identification.

Definition 2.6 (Bayesian identification) *The Bayesian model \mathcal{E}_b is Bayesian identified if and only if $\mathcal{A}\mathcal{X} = \mathcal{A}$.*

More generally, the concept of identification may be introduced for a sub- σ -field (*i.e.* a parameter) $\mathcal{B} \subset \mathcal{A}$:

Definition 2.7 (Bayesian identified parameter) *\mathcal{B} is an \mathcal{E}_b -identified parameter if $\mathcal{B}\mathcal{S} = \mathcal{B}$.*

We recall that $\mathcal{B}\mathcal{S}$ denotes the projection of \mathcal{S} on \mathcal{B} , that is,

$$\mathcal{B}\mathcal{S} := \sigma(\{\mathbf{E}[s|\mathcal{B}]; s \text{ belongs to the set of positive random variables defined on } (X, \mathcal{X})\}),$$

where $\sigma(\{G\})$ denotes the σ -field generated by the set G .

Definition 2.6 means that in a Bayesian identified model, \mathcal{A}_*^μ is almost surely equal to \mathcal{A} , that is, $\overline{\mathcal{A}_*} \cap \mathcal{A} = \mathcal{A}$ or: $\forall A \in \mathcal{A}, \exists B \in \mathcal{A}_*$ such that $\mu(A \triangle B) = 0$, where \triangle denotes the symmetric difference. This property shows that Bayesian identification is an almost sure measurable identification with respect to the prior. We also have the following characterization of Bayesian identification, see Florens et al. (1990, Theorem 4.6.21).

Theorem 2.2 *If (Θ, \mathcal{A}) is a Souslin space (see definition before Theorem 2.1) and if \mathcal{X} is separable, then the Bayesian model is identified if and only if $\exists \Theta_0 \subset \mathcal{A}$ such that (i) $\mu(\Theta_0) = 0$ and (ii) the sampling model restricted to $\Theta - \Theta_0$ is identified, *i.e.* $\theta \rightarrow P^\theta$ is injective on $\Theta - \Theta_0$.*

It is clear from the theorem that sampling identification (and equivalently measurable identification) implies Bayesian identification but the reverse is not true. For instance, a degenerate prior that puts mass

one on specific values of the parameter space makes Bayesian identified a model that is not sampling identified as the following trivial example shows. Suppose that we observe a random variable x from the sampling model $x|\theta_1, \theta_2 \sim \mathcal{N}(\theta_1 + \theta_2, 1)$, $\theta = (\theta_1, \theta_2)' \in \Theta = \mathbb{R}^2$, and that we endow the parameter θ with the degenerate prior $\theta \sim \mathcal{N}((\theta_{10}, \theta_{20})', \iota')$, where $\iota = (1, 1)'$. This prior gives probability zero to all the value in \mathbb{R}^2 but the values on the line $\theta_2 = \theta_1$. Therefore, $\Theta_0 = \mathbb{R}^2 \setminus \{\theta \in \mathbb{R}^2; \theta_2 = \theta_1\}$ and the sampling distribution is injective on $\Theta - \Theta_0 = \{\theta \in \mathbb{R}^2; \theta_2 = \theta_1\}$. This model is not sampling identified nor measurable identified. However, this specification of the prior makes the model Bayesian identified.

Before concluding this section we point out the following relationships existing between the three concepts of identification that we have seen. (i) Measurable identification implies Bayesian identification for any μ . (ii) Measurable identification implies sampling identification if and only if \mathcal{A} is separating, that is, all the atoms of \mathcal{A} are singletons.

2.4 Identification and Bayesian consistency

In this section we present the important connection that exists between the concepts of identification and of *exact estimability* in the Bayesian models. We start by defining exact estimability.

Definition 2.8 (Exact estimability) *In a Bayesian model \mathcal{E}_b , let \mathcal{B} be a sub- σ -field of \mathcal{A} . Then, \mathcal{B} is exactly estimable if and only if $\mathcal{B} \subset \overline{\mathcal{X}}$ where $\overline{\mathcal{X}}$ is the σ -field generated by \mathcal{X} and all the null sets of the product space $\Theta \times X$ with respect to Π . Equivalently, \mathcal{B} is exactly estimable if $\mathcal{B} \perp \mathcal{A}|\overline{\mathcal{X}}$.*

The σ -field $\overline{\mathcal{X}}$ is a σ -field on the product space $\Theta \times X$ and not a σ -field on the sampling space X . So, it is possible that $\mathcal{B} \subset \overline{\mathcal{X}}$ but that $\mathcal{B} \not\subset \mathcal{X}$. For instance, consider $x_i|\theta \stackrel{iid}{\sim} \mathcal{N}(\theta, 1)$ and $\mathcal{B} = \sigma(\{\theta\})$. Consider a prior for θ that puts all its mass on the sample mean \bar{x} , then we clearly have that $\mathcal{B} \subset \overline{\mathcal{X}}$ but $\mathcal{B} \not\subset \mathcal{X}$. This situation is artificial in small sample, but it describes well what it happens asymptotically if \bar{x} is a consistent estimator since $\bar{x} \rightarrow \theta$, Π -a.s. without requiring a degenerate prior, see also Definition 2.9 and the discussion below it.

The inclusion $\mathcal{B} \subset \overline{\mathcal{X}}$ means that for any positive random variable a defined on (Θ, \mathcal{B}) , the posterior expectation $\mathbf{E}(a|\mathcal{X}) = a$ Π -a.s., provided that the conditional expectation exists. This means that

a posteriori – that is, after observing the sample – we know a Π -a.s. Thus, any sub- σ -field of $\mathcal{A} \cap \overline{\mathcal{X}}$ is exactly estimable. The next proposition states a link between exact estimability and Bayesian identification.

Proposition 2.2 *Any exactly estimable parameter $\mathcal{B} \subset \mathcal{A} \cap \overline{\mathcal{X}}$ is Bayesian identified.*

The result in the proposition, together with the following result, allows to show that in an *i.i.d.* experiment, the minimal sufficient σ -field is Bayesian identified.

Theorem 2.3 *In an *i.i.d.* model, the minimal sufficient σ -field of \mathcal{A} , $\mathcal{A}\mathcal{X}$, is exactly estimable.*

Moreover, if $\mathcal{A}\mathcal{X}$ is exactly estimable, then the following theorem shows that the reverse implication of Proposition 2.2 holds.

Theorem 2.4 *Let $\{(\Theta, \mathcal{A}), (X, \mathcal{X}), \Pi\}$ be a Bayesian model such that the Bayesian minimal sufficient σ -field $\mathcal{A}\mathcal{X}$ is exactly estimable, i.e. $\mathcal{A}\mathcal{X} \subset \overline{\mathcal{X}}$. Then, $\mathcal{B} \subset \mathcal{A}$ is exactly estimable if and only if \mathcal{B} is *a.s. identified*, i.e. $\mathcal{B} \subset \mathcal{A}\mathcal{X}$.*

The concept of exact estimability has a particular interest in asymptotic models. Let $x^{(n)} := (x_1, \dots, x_n)$ denote a sequence of n observations and $x^{(\infty)}$ denote the sample of infinite size. The σ -field generated by $x^{(n)}$ (resp. $x^{(\infty)}$) is denoted by \mathcal{X}_n (resp. \mathcal{X}_∞). For any function $a(\cdot)$ of the parameter θ , the martingale convergence theorem implies that $\mathbf{E}(a|\mathcal{X}_n) \rightarrow \mathbf{E}(a|\mathcal{X}_\infty)$ a.s. This convergence is a.s. with respect to the joint distribution Π and also with respect to the predictive probability P . This is one concept of *Bayesian consistency* for which the convergence must be taken with respect to the joint probability distribution Π . Another concept of Bayesian consistency is convergence with respect to the sampling measure P^θ , that is, $\mathbf{E}(a|\mathcal{X}_n) \rightarrow a(\theta)$, P^θ - a.s. This convergence does not follow from the previous argument and there are cases in which it is not be verified, see *e.g.* Diaconis and Freedman (1986) and Florens and Simoni (2012).

From the above concept of Bayesian consistency, if $\mathcal{B} \subset \overline{\mathcal{X}_\infty}$, we have $\mathbf{E}(a|\mathcal{X}_n) \rightarrow a$ Π -a.s. for any \mathcal{B} -measurable positive function a of θ , that is, the posterior mean of a converges a.s. to the true model. This is a consequence of the asymptotically exact estimability which we now define.

Definition 2.9 (Asymptotically exact estimability) Let \mathcal{E}_∞ be the sequential Bayesian experiment defined as $\mathcal{E}_\infty = (\Theta \times X, \mathcal{A} \otimes \mathcal{X}, \Pi, \mathcal{X}_n \rightarrow \mathcal{X}_\infty)$ and \mathcal{B} be a sub- σ -field of \mathcal{A} such that $\mathcal{B} \subset \overline{\mathcal{X}}_\infty$. Then, \mathcal{B} is said to be asymptotically exactly estimable in \mathcal{E}_∞ .

This definition and the martingale convergence theorem imply that if $\mathcal{B} \subset \mathcal{A}$ is asymptotically exactly estimable then, for every integrable real random variable a defined on \mathcal{B} , $\mathbf{E}(a|\mathcal{X}_n) \rightarrow a$ Π -a.s.

In a Souslin space, *asymptotic exact estimability* means that the posterior distribution is asymptotically a Dirac measure on a function of the sample, see Florens et al. (1990, Theorem 4.8.3). The intuition is that, if $\mathbf{E}(a|\mathcal{X}_n) \rightarrow a$ Π -a.s., then also $\mathbf{E}(a^2|\mathcal{X}_n) \rightarrow a^2$ Π -a.s. which implies that the posterior variance of a converges to 0. This clearly explains that the posterior distribution concentrates around the Π -a.s. limit of its mean. More generally, an integrable real random variable $a(\cdot)$ defined on \mathcal{B} is asymptotically exactly estimable if there exists a strongly consistent sequence of estimators of $a(\theta)$, i.e. if there exists a set of random variables $(t_n)_{n \in \mathbb{N}}$, each defined on \mathcal{X} , such that $t_n \rightarrow a(\theta)$, Π -a.s. In sampling theory, a necessary condition for the existence of such a sequence is the identifiability of the parameter.

2.5 Nonidentified and partially identified models

Between the two extreme situations of identification and non-identification there are models that are *partially identified*. Partial identification arises when the combination of available data and assumptions that are plausible for the model only allows to place the population parameter of interest γ within a proper subset Γ_I of the parameter space Γ called identified set. If one has very precise prior information, identification can be restored by specifying a prior distribution degenerated at some point in Γ_I as we explain in section 3.1 - Model 5. However, this strategy can be unsuitable if one is concerned with robustness of the Bayesian procedure. Therefore, a more appealing approach consists in working directly with the parameter Γ_I which is a set and: first, define the prior and posterior distribution of Γ_I (see section 4); second, define a prior for γ inside this set (see section 5.2).

In a partially identified model, usually, there is an identified parameter, say θ , and a partially identified parameter which is the parameter of interest, denoted by γ . The latter enters the model as a supplementary parameter and is linked to θ through a relation of the form $A(\theta, \gamma) \in A_0 \subset \Phi$ where Φ

is a suitably defined space and A is a given function of (θ, γ) . Such relation characterizes the identified set as $\Gamma_I := \{\gamma; A(\theta, \gamma) \in A_0 \in \Phi\}$ which may not be a singleton. The function $(\theta, \gamma) \mapsto A(\theta, \gamma)$ may be a parametric likelihood function, flat over a region of γ s, and A_0 is its maximum value. Alternatively, Γ_I may be a set of moment restrictions as in **Example 1** below. See Chen et al. (2018) for more details on Bayesian analysis of these examples.

Example 1 (Moment Restrictions) *A natural construction of sampling econometric model which may lead to identification issues is the following. The econometrician first considers a sampling model perfectly identified $\{\Theta, (X, \mathcal{X}), (P^\theta)_{\theta \in \Theta}\}$ but she completes her specification by considering another parameter $\gamma \in \Gamma$. This parameter γ is the natural parameter arising in economic models and is related to θ . The relation linking θ and γ may be of the form $A(\theta, \gamma) = 0$ or $A(\theta, \gamma) \geq 0$ or more generally $A(\theta, \gamma) \in A_0$, $A_0 \in \Phi$ where Φ is a suitably defined space and A is a given function of (θ, γ) . The parameter γ may be identified or not depending on the relation A . For example, the relation A may be a set of moment restrictions specified with respect to the distribution F which generates the data. In that case $\theta = F$ the parameter γ could be characterized through the moment condition*

$$A(\theta, \gamma) := \mathbf{E}_F(h(x, \gamma)) = 0, \quad (2.1)$$

where h is a known moment function and $\mathbf{E}_F(\cdot)$ denotes the expectation taken with respect to F . The parameter γ is identified if a unique solution to equation (2.1) exists. Condition (2.1) may be extended into $\mathbf{E}_F(h(x, \gamma)) \geq 0$ which in general defines a set of γ solutions to these inequalities for a given F .

To develop Bayesian analysis it is important to understand that the specification of the prior distribution differs in nonidentified models and in partially identified models. In partially identified models, the prior distribution is naturally decomposed in the marginal prior on θ and the conditional prior on γ given θ which incorporates the link between θ and γ : $x|\theta, \gamma \sim P^\theta$, $\theta \sim \mu_\theta$ and $\gamma|\theta \sim \mu_\gamma(\cdot|\theta)$.

On the other hand, models with nonidentified parameters arise when θ is partitioned in $\theta = (\beta, \gamma)$ where β is identified and γ is not. For instance, γ can be the parameter of the distribution of the heterogeneity parameter β . Therefore, the prior is naturally decomposed in the prior on β conditional on γ , $\mu_\beta(\cdot|\gamma)$, and the prior μ_γ on γ . This decomposition of the prior distribution is the reverse of the decomposition made for partially identified models.

3 Bayesian analysis of nonidentified models

This section is made of two subsections that describe two different options one has to make Bayesian inference for nonidentified models. We first discuss and construct models that may satisfy Bayesian identification even if they are not sampling identified nor measurable identified. This happens when the function $\theta \mapsto P^\theta$ is not injective but the prior assigns a mass zero on the nonidentified component of the model, as it has been shown in Theorem 2.2. A second option, discussed in section 3.2, consists in working with the nonidentified model without introducing any identifying priors or identifying assumptions. The beauty of the Bayesian approach is that, as long as the prior is proper, one can make Bayesian inference because the posterior distribution is well defined and then MCMC algorithms can be designed to simulate from the posterior even in a nonidentified model.

3.1 Models with heterogeneity and hyperparameters

This section considers models where the parameter of interest γ is a parameter of the distribution of a heterogeneity parameter β . More precisely, the parameter space (Θ, \mathcal{A}) contains two subparameter spaces B and Γ with the associated sub- σ -fields \mathcal{B} and \mathcal{G} , where \mathcal{B} is the μ -a.s. minimal sufficient σ -field in the parameter space. The σ -field \mathcal{G} is not identified in the sense that $\mathcal{G} \not\subseteq \mathcal{B}$. The sampling distribution only depends on \mathcal{B} and the Lebesgue data density, if it exists, satisfies $p(x|\beta, \gamma) = p(x|\beta)$ a.s., where $\beta \in (B, \mathcal{B})$ and $\gamma \in (\Gamma, \mathcal{G})$. The posterior satisfies $\mu(\gamma|x, \beta) = \mu(\gamma|\beta)$ a.s. An interesting situation, known as local identification, arises when the σ -field \mathcal{G} is not identified only for particular values taken on by the parameter \mathcal{B} , that is, $p(x|\beta = \bar{\beta}, \gamma) = p(x|\beta = \bar{\beta})$ for some value $\bar{\beta}$, see *e.g.* Drèze and Richard (1983), Kleibergen and van Dijk (1998), Hoogerheide et al. (2007) and Kleibergen and Mavroeidis (2014). We do not analyze this situation in this paper.

The prior is naturally specified as the product of the conditional distribution of \mathcal{B} given \mathcal{G} and the marginal on \mathcal{G} . In this case, \mathcal{G} is interpreted as a parameter of the prior on \mathcal{B} , usually called an *hyperparameter*, see *e.g.* Berger (1985), or *latent variable*. Even if this model is not sampling (nor measurable) identified, it may satisfy Bayesian identification if the prior is conveniently specified as it has been shown in Theorem 2.2. This identification by the prior is fully artificial in the finite-dimensional parameter case where a degenerate prior is required to get identification and we do not

consider this case. Instead, we consider nonparametric and asymptotic settings where identification can be obtained without a degenerate prior. The main argument that we use to get Bayesian identification in this way is the following.

Proposition 3.1 *Let us consider the Bayesian model $\mathcal{E}_b = \{\Theta \times X, \mathcal{A} \otimes \mathcal{X}, \Pi\}$ and a sub- σ -field $\mathcal{B} \subset \mathcal{A}$ such that $\mathcal{B} = \mathcal{A}\mathcal{X}$. Then, any sub- σ -field $\mathcal{G} \subset \mathcal{A}$ is identified if and only if $\mathcal{G} \subset \mathcal{B}$, μ -a.s. or, equivalently, if any \mathcal{G} -measurable function is μ -a.s. equal to a \mathcal{B} -measurable function.*

Notice that the *a.s.* in the proposition are related to the prior and that \mathcal{B} is the minimal sufficient σ -field. The proposition states that any sub- σ -field $\mathcal{G} \subset \mathcal{A}$ is identified if and only if it is almost surely included in \mathcal{B} . Intuitively, the result of the proposition holds if the conditional prior distribution on \mathcal{G} given \mathcal{B} is degenerate into a Dirac measure on a \mathcal{B} -measurable function. No requirement are made on the marginal prior on \mathcal{G} . We now discuss four models where Bayesian identification can be obtained without a degenerate marginal prior on γ and a fifth example where the model is partially identified, and contrarily to the previous cases, Bayesian identification can be only obtained artificially with a degenerate marginal prior.

Model 1: unobserved heterogeneity (or incidental parameter). The incidental parameter problem typically arises with panel data models when a regression model includes an agent specific intercept β_i which is a latent variable. The general conditional model can be simplified as

$$x_{it}|\beta_i, \gamma \sim \text{ind } \mathcal{N}(\beta_i, \sigma^2), \quad i = 1, \dots, \infty, \quad t = 1, \dots, T$$

where σ^2 is known and γ is a hyperparameter characterizing the distribution of β_i . If β_i is *i.i.d.* across individuals then γ is the common parameter in the population. Let $\beta := (\beta_1, \beta_2, \dots)$. The μ -a.s. minimal sufficient σ -field \mathcal{B} is generated by β , that is, β is the identified parameter. Instead of estimating the distribution of the heterogeneity parameter β_i one could be satisfied with the estimation of the common parameter γ . Unfortunately, γ is not measurably identified even if it may be Bayesian identified. To illustrate this, let us specify a prior probability for β and γ as follows:

$$\beta_i|\gamma \sim \text{i.i.d. } \mathcal{N}(\gamma, \sigma_0^2), \quad i = 1, \dots, \infty$$

$$\gamma \sim \mu_\gamma,$$

where σ_0^2 is known and μ_γ is a probability measure. By the strong law of large numbers, $\gamma = \lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \beta_i$ *a.s.* with respect to both the conditional prior distribution of β given γ and the joint prior distribution of (β, γ) . Therefore, the conditional distribution of $\gamma | \{\beta_i\}_{i=1}^\infty$ puts all its mass on $\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \beta_i$ and the model is fully Bayesian identified. We stress that this property does not depend on the marginal prior on γ , which does not have to be degenerate on some value, but crucially depends on the prior of β given γ .

Model 2: Gaussian process and hyperparameter in the mean. Let $L^2[0, 1]$ denote the space of square integrable functions defined on $[0, 1]$ with respect to the uniform distribution, endowed with the L^2 -inner product $\langle \cdot, \cdot \rangle$ and the induced norm $\| \cdot \|$. We consider a sample space $X = L^2[0, 1]$. The sample is made of a single observation which is a trajectory x from a Gaussian process in X : $x | \beta, \gamma \sim \mathcal{GP}(\beta, \Sigma)$ where $\beta \in X$ and $\Sigma : X \rightarrow X$ is a covariance operator which is bounded, linear, positive definite, self-adjoint and trace-class. It follows that $\mathbf{E} \|x\|^2 < \infty$. The parameter β is measurably identified but the parameter γ is not and it may be interpreted as a low dimensional parameter that controls the distribution of β .

Suppose that β and γ are two random variables with values in (B, \mathcal{B}) and $(\mathbb{R}_+, \mathcal{R})$, respectively, with $B = L^2[0, 1]$ and \mathcal{B} and \mathcal{R} are the associated Borel σ -field. The joint prior distribution on $(B \times \mathbb{R}_+, \mathcal{B} \otimes \mathcal{R})$ is specified as:

$$\begin{aligned} \beta | \gamma &\sim \mathcal{GP}(\gamma \beta_0, \sigma_0^2 \Omega), & \beta_0 \in B = L^2[0, 1], & \quad \Omega : B \rightarrow B \\ \gamma &\sim \mathcal{N}(\gamma_0, \sigma_0^2), & \gamma_0, \sigma_0^2 \in \mathbb{R}_+, & \end{aligned}$$

where β_0 , γ_0 , σ_0^2 and Ω are known hyperparameters. The covariance operator Ω is bounded, linear, positive definite, self-adjoint and trace-class and its eigensystem is denoted by $(\lambda_j, \varphi_j)_{j \geq 1}$. By definition of Gaussian processes in Hilbert space, $\beta | \gamma \sim \mathcal{GP}(\gamma \beta_0, \sigma_0^2 \Omega)$ if and only if $\langle \beta, \varphi_j \rangle | \gamma \sim$

$\mathcal{N}(\langle \beta_0, \varphi_j \rangle \gamma, \sigma_0^2 \lambda_j)$ for all $j \geq 1$. Hence, it follows that

$$\frac{\langle \beta, \varphi_j \rangle}{\sigma_0 \sqrt{\lambda_j}} \Big| \gamma \sim \mathcal{N}\left(\langle \beta_0, \varphi_j \rangle / (\sigma_0 \sqrt{\lambda_j}) \gamma, 1\right).$$

Let us write $\beta = \beta_0 \gamma + U$, where $U \sim \mathcal{GP}(0, \sigma_0^2 \Omega)$. Let $\gamma^*(\beta)$ denote the minimizer of the least squares criterion: $\gamma^*(\beta) := \arg \min_{\gamma} \sum_{j=1}^{\infty} (\langle \beta, \varphi_j \rangle - \langle \beta_0, \varphi_j \rangle \gamma)^2 / (\sigma_0^2 \lambda_j)$ which takes the form

$$\gamma^*(\beta) = \frac{\sum_{j=1}^{\infty} \frac{\langle \beta, \varphi_j \rangle \langle \beta_0, \varphi_j \rangle}{\sigma_0^2 \lambda_j}}{\sum_{j=1}^{\infty} \frac{\langle \beta_0, \varphi_j \rangle^2}{\sigma_0^2 \lambda_j}} = \gamma + \frac{\sum_{j=1}^{\infty} \xi_j \frac{\langle \beta_0, \varphi_j \rangle}{\sqrt{\lambda_j}}}{\sum_{j=1}^{\infty} \frac{\langle \beta_0, \varphi_j \rangle^2}{\sigma_0 \lambda_j}}, \quad (3.1)$$

where $\xi_j := \langle U, \varphi_j \rangle / (\sigma_0 \sqrt{\lambda_j}) \sim i.i.d. \mathcal{N}(0, 1)$, $j \geq 1$. If $\langle \Omega^{-1/2} U, \Omega^{-1/2} \beta_0 \rangle = 0$ then the second term on the right hand side of (3.1) is zero and so $\gamma^*(\beta) = \gamma$. It follows that we can write γ as a function of β and we have exact estimability and Bayesian identification of γ , that is, $\gamma \in \bar{\mathcal{B}}$. As in **Model 1**, this property does not depend on the marginal prior on γ but crucially depends on the prior of β given γ . The condition $\langle \Omega^{-1/2} U, \Omega^{-1/2} \beta_0 \rangle = 0$ means that the scaled error term and the prior mean functions have to be orthogonal in $L^2[0, 1]$.

Model 3: Gaussian process and hyperparameter in the variance. Let us consider the same setting as in Model 2 but suppose that now we are interested in the common variance parameter σ^2 . That is, the prior distribution for $(\beta, \sigma^2) \in (B \times \mathbb{R}_+, \mathcal{B} \otimes \mathcal{R})$ is specified as:

$$\beta | \sigma^2 \sim \mathcal{GP}(0, \sigma^2 \Omega), \quad \text{and} \quad \sigma^2 \sim \mathcal{IG}(\nu_0, \sigma_0^2),$$

where \mathcal{IG} denotes an inverse-gamma distribution and $\Omega : X \rightarrow X$ is a known covariance operator which is bounded, linear, positive definite, self-adjoint and trace-class. The hyperparameters ν_0 and σ_0^2 are known. Knowledge of Ω implies knowledge of its eigensystem $(\lambda_j, \varphi_j)_{j \geq 1}$.

By definition, $\beta | \sigma^2 \sim \mathcal{GP}(0, \sigma^2 \Omega)$ if and only if $\left\{ \langle \beta, \varphi_j \rangle / \sqrt{\lambda_j} \right\}_{j \geq 1} | \sigma^2 \sim i.i.d. \mathcal{N}(0, \sigma^2)$. By the strong law of large numbers we have

$$\lim_{J \rightarrow \infty} \frac{1}{J} \sum_{j=1}^J \frac{\langle \beta, \varphi_j \rangle^2}{\lambda_j} = \sigma^2, \quad \mu - a.s.$$

This convergence is valid with respect to both the conditional prior $\mu(\beta|\sigma^2)$ of β given σ^2 and the joint prior $\mu(\beta, \sigma^2)$ of (β, σ^2) and shows that σ^2 is a $\mu - a.s.$ function of β . Therefore, σ^2 is Bayesian identified.

Model 4: Bayesian nonparametric and Dirichlet process. We observe an n -sample from the following sampling model $x_i|F, G \sim i.i.d. F, i = 1, \dots, n$, where F and G are two probability distributions – for instance on \mathbb{R} . We can interpret F as a nuisance parameter and G is the only parameter of interest which characterizes the distribution of F . The parameter F generates a σ -field \mathcal{B} that is measurable identified while G generates a σ -field \mathcal{G} that is not measurable identified. We specify a nonparametric Dirichlet process prior for F with parameters n_0 and G : $F|G \sim Dir(n_0, G)$ and $G \sim \mu$, where μ denotes a probability measure on the space of distributions that generate almost surely a diffuse distribution, *i.e.* $G(x) = 0, \forall x$. The F generated from the above Dirichlet process can be equivalently generated by using the stick-breaking representation as $F = \sum_j \alpha_j \delta_{\xi_j}$, where $\{\xi_j\}_{j \geq 1}$ are independent draws from G , $\alpha_j = v_j \prod_{k=1}^j (1 - v_k)$ with $\{v_j\}_{j \geq 1}$ independent draws from a Beta distribution $Be(1, n_0)$ and $\{v_j\}_{j \geq 1} \perp \{\xi_j\}_{j \geq 1}$, see Appendix D in the Supplement. Then, $\lim_{J \rightarrow \infty} \frac{1}{J} \sum_{j=1}^J \delta_{\xi_j} = G, \mu - a.s.$ which proves that G is a μ -a.s. function of F and so, it is Bayesian identified.

Model 5: Moment conditions and partially identified models. Consider the setting of **Example 1** where the parameter of interest $\gamma \in (\Gamma, \mathcal{G})$ is characterized through a moment condition of the type $\mathbf{E}_F(h(x, \gamma)) \geq 0$ for a known function h and a data generating process F . We observe n realisations (x_1, \dots, x_n) of n *i.i.d.* random variables from F and specify for F a Dirichlet process prior with parameters n_0 and F_0 :

$$\begin{aligned} x_i|F &\sim i.i.d. F, & i = 1, \dots, n, \\ F &\sim Dir(n_0, F_0). \end{aligned}$$

The parameter F generates the σ -field \mathcal{B} and is measurable identified. Suppose that γ is partially identified. Despite of this, γ can be Bayesian identified by specifying a conditional prior $\mu(\gamma|F)$ for γ , given F , degenerate on a given functional of F . For example, if the restriction $\mathbf{E}_F(h(x, \gamma)) \geq 0$

writes $\gamma \in [\phi_1(F), \phi_2(F)]$ for two functionals ϕ_1, ϕ_2 of F , then $\mu(\gamma|F)$ could be specified as a Dirac on $(\phi_1(F) + \phi_2(F))/2$. We stress that Bayesian identification in this model is obtained artificially and in a different way than in Models 1-4 above because of the construction of a conditional degenerate prior for the partially identified parameter γ , given the identified one. Instead, in Models 1-4 we have specified a marginal prior for the unidentified parameter that was not degenerate and Bayesian identification arose because of the specification of the conditional prior for the identified parameter.

This artificial way to obtain Bayesian identification through a degenerate prior can be reprehensible as it can be seen against the logic of partial identification. In section 4 we will proceed without imposing Bayesian identification.

3.2 Latent variable models

In this section we take the example of latent variable models to illustrate that it is possible to make Bayesian inference directly for the nonidentified model without introducing an identifying prior or identifying assumptions. Because these models are well known in the literature we discuss them briefly.

Consider discrete observable outcomes y_i arising from the model $y_i = g(z_i)$, where $g(\cdot)$ is a given function and z_i is a latent variable satisfying: $z_i = x_i'\delta + u_i$, where x_i is an observable real-valued random vector, δ is the parameter vector of interest, and u_i is an unobservable component with unrestricted variance. Identification problems appear if the $g(\cdot)$ function is invariant to location or scale transformations of z . Differently from frequentist analysis, imposing identifying restrictions is not necessary if one wants to conduct Bayesian analysis. With proper priors, one obtains the posterior for the nonidentified model, constructs an MCMC algorithm to simulate from this posterior and then marginalises to obtain the posterior of the identified parameters. It is often the case that it is easier to construct MCMC algorithms for unrestricted parameters and they have better mixing properties than MCMC algorithms for the corresponding identified model, see *e.g.* van Dyk and Meng (2001).

The simpler example of a latent variable model is the binary logit / probit, where the observable outcome y_i is binary and modeled as: $y_i = \mathbb{1}\{z_i > 0\}$ with z_i a real-valued random variable. If u_i follows a Gaussian (resp. Logistic) distribution we have the probit (resp. logit) model. The identification problem arises because the indicator function is invariant to scale transformation of z_i : multiplication

of z_i by a positive constant does not change the likelihood. Identification can be restored by imposing for instance the restriction $Var(u_i) = 1$. However, Bayesian analysis does not require any identifying restrictions. Other example of latent variable models are provided in Appendix A.

4 Bayesian analysis of partially identified models

This section considers partially identified models as described in section 2.5. Let us consider a model with an identified parameter θ and another parameter γ characterized by the condition

$$A(\theta, \gamma) \in A_0, \tag{4.1}$$

where $\gamma \in \Gamma \subset \mathbb{R}^{d_\gamma}$, $A : (\Theta, \mathcal{A}, \mu) \times \Gamma \rightarrow \Phi$, for some finite or infinite dimensional space Φ and A_0 is a subset of Φ . The function $A(\theta, \gamma)$ is usually either a likelihood function or a moment function and condition (4.1) can contain both inequalities and equalities, see Chen et al. (2018) for interesting examples. The parameter γ is the parameter of interest and, depending on the relation (4.1), it may be only partially identified which means that for a given θ there may exist more than one value of γ in Γ satisfying (4.1) but not all the values in Γ satisfy the condition.

Let $\Gamma_I := \Gamma_I(\theta) = \{\gamma \in \Gamma; A(\theta, \gamma) \in A_0\}$ be the *identified set* which is a proper subset of Γ and depends on θ . The model is point identified if Γ_I is a singleton and is partially identified otherwise. Inference of partially identified models can focus either on γ , or on the set Γ_I or on both. In this section, we focus on Γ_I and so do not endow γ with a prior. A prior on γ will be introduced in section 5.

In the Bayesian analysis, the quantity $A(\theta, \gamma)$ in (4.1) is a random element, where the randomness comes from θ : for each $\gamma \in \Gamma$, $A(\cdot, \gamma)$ is a function on (Θ, \mathcal{A}) which takes values in the state space $(\Phi, \mathcal{B}(\Phi))$, where $\mathcal{B}(\Phi)$ denotes the Borel σ -field associated with Φ . Therefore, condition (4.1) has to hold *a.s.* with respect to the prior distribution of θ . For a realisation of θ in Θ , the set Γ_I is constructed on the basis of the trajectory $\gamma \mapsto A(\theta, \gamma)$, $\gamma \in \Gamma$. Therefore, $\Gamma_I(\theta)$ has to be interpreted as a random set where the randomness comes from θ . This differs from the frequentist analysis of partially identified models where Γ_I is non random.

The next definition describes a random closed set. For this, let \mathcal{C} (resp. \mathcal{O}) denote a family of

closed (resp. open) subsets of Γ .

Definition 4.1 (Random closed set) *Let $(\Theta, \mathcal{A}, \mu)$ be a complete probability space and Γ be a locally compact space. The multivalued function $\Gamma_I(\cdot) : \Theta \rightarrow \mathcal{C}$ is called a random closed set if, for every compact set K in Γ , $\{\theta; \Gamma_I(\theta) \cap K \neq \emptyset\} \in \mathcal{A}$.*

The multivalued function $\Gamma_I(\cdot) : \Theta \rightarrow \mathcal{O}$ is called a *random open set* if its complement is a random closed set. The following proposition gives a condition that guarantees that Γ_I is a random closed set. A regular closed set C is such that C coincides with the closure of its interior, see Molchanov (2005).

Proposition 4.1 *Assume that the multivalued function $\Gamma_I(\cdot) : \Theta \rightarrow \mathcal{C}$ has realizations a.s. equal to regular closed subsets of Γ . If the set A_0 in (4.1) is such that $A_0 \in \mathcal{B}(\Phi)$ then $\Gamma_I(\cdot)$ is a closed random set.*

We assume in the following that $\Gamma_I(\theta)$ is a closed random set. Our analysis can be extended to the case where $\Gamma_I(\theta)$ is an open random set. It is convenient to introduce the stochastic process $\{g_\theta(\gamma)\}_{\gamma \in \Gamma}$ associated with condition (4.1), where for every $\gamma \in \Gamma$

$$g_\theta(\gamma) := \mathbb{1}\{A(\theta, \gamma) \in A_0 \subset \Phi\}. \quad (4.2)$$

So, $g_\theta(\gamma)$ is the indicator of the random set $\Gamma_I(\theta) \subset \Gamma$ and γ is the index of the stochastic process $g_{(\cdot)}(\cdot) : (\Theta, \mathcal{A}, \mu) \times \Gamma \rightarrow \{0, 1\}$. When $\Gamma_I(\theta)$ is a separable random set, see Molchanov (2005, Definition 4.8), then it can be completely described through its indicator function.

4.1 Construction of the prior and posterior capacity functionals

A Bayesian analysis requires the specification of a prior distribution for the random set Γ_I . Since Γ_I depends on θ , we propose to construct such a prior by first specifying a prior distribution for θ and then recovering from it the prior for Γ_I . Suppose $X \subseteq \mathbb{R}$ and let F denote the data distribution. We suppose that θ may be written as a measurable functional of the distribution F , that is, there exists a measurable functional $\phi(\cdot)$ such that $\theta = \phi(F)$. For instance, $\theta = \mathbf{E}_F(x) = \int x dF(x)$.

In our analysis F is not restricted to belong to some parametric class. Thus, we specify a Dirichlet

process prior for F and write $F \sim \text{Dir}(n_0, F_0)$ where $n_0 \in \mathbb{R}_+$ and F_0 is a diffuse probability measure on \mathbb{R} , i.e. $F_0(y) = 0, \forall y \in \mathbb{R}$, see e.g. Ferguson (1973) and Florens (2002). The prior distribution for $\Gamma_I(\theta)$ is obtained from this prior for F through the *prior capacity functional*. Define \mathcal{K} as the family of compact subsets of Γ . The *prior capacity functional* $T_{\Gamma_I} : \mathcal{K} \mapsto [0, 1]$ is given by

$$T_{\Gamma_I}(K) := P\{K \cap \Gamma_I(\theta) \neq \emptyset\}, \quad K \in \mathcal{K},$$

where the probability P is determined by the prior distribution of θ which in turns is determined by the Dirichlet process prior on F : $F \sim \text{Dir}(n_0, F_0)$ since $\theta = \phi(F)$.

When the prior capacity functional is defined on singletons instead of on \mathcal{K} , i.e. $K = \{\gamma\}$ for some $\gamma \in \Gamma$, then it is called *prior coverage function* of Γ_I and denoted by $p_{\Gamma_I}(\gamma)$. In particular, $\forall \gamma \in \Gamma$

$$p_{\Gamma_I}(\gamma) := \mathbf{E}(g_\theta(\gamma)) = \text{Prob}(g_\theta(\gamma) = 1) = P(\{\gamma\} \cap \Gamma_I(\theta) \neq \emptyset) = P(A(\theta, \gamma) \in A_0) \in [0, 1], \quad (4.3)$$

where $g_\theta(\cdot)$ is the stochastic process defined in (4.2) and P and \mathbf{E} are the probability and expectation, respectively, taken with respect to the prior of θ . The appealing fact of the prior coverage function with respect to the prior capacity functional is that $p_{\Gamma_I}(\gamma)$ can be represented graphically in an easier way than $T_{\Gamma_I}(K)$ (at least if $\Gamma \subset \mathbb{R}$), see for instance Figures 1, 3 and 9 where we represent $p_{\Gamma_I}(\cdot)$ and the posterior coverage function $p_{\Gamma_I}(\cdot|x)$. We remark that the prior coverage function $p_{\Gamma_I}(\gamma)$ does not characterize the distribution of the stochastic process $g_\theta(\gamma)$ which is instead characterized by its finite-dimensional distributions.

We detail now the construction of the prior and posterior capacity functionals with the help of a generic example. Suppose $\theta := (\theta_1, \theta_2)' \in \mathbb{R}^2$, $A(\theta, \gamma) = (\theta_1 - \gamma, \gamma - \theta_2)'$ for some $\gamma \in \mathbb{R}$ and $A_0 = (-\infty, 0] \times (-\infty, 0]$. Hence, the condition $A(\theta, \gamma) \in A_0$ writes as $\gamma \in [\theta_1, \theta_2]$ and $\Gamma_I = [\theta_1, \theta_2]$. We recall that the aim is to make inference on the identified set Γ_I and not on the partially identified parameter. To start with, suppose that we specify a parametric prior for θ . For instance, $\theta_1 \sim \mathcal{U}[0, 1]$ and $\theta_2 \sim \mathcal{U}[1, 2]$. For any $K \in \mathcal{K}$, write $K = [K_1, K_2]$, $\bar{K} = [\bar{K}_1, \bar{K}_2] := K \cap [0, 1]$ and $\bar{\bar{K}} =$

$[\bar{K}_1, \bar{K}_2] := K \cap [1, 2]$. The prior capacity functional then is:

$$T_{\Gamma_I}(K) = \begin{cases} 0 & \text{if } K \cap [0, 2] = \emptyset \\ \bar{K}_2 - \bar{K}_1 & \text{if } K \cap [0, 1] \neq \emptyset \text{ and } K \cap [1, 2] = \emptyset \\ \bar{K}_2 - \bar{K}_1 + \bar{\bar{K}}_2 - \bar{\bar{K}}_1 & \text{if } K \cap [0, 1] \neq \emptyset \text{ and } K \cap [1, 2] \neq \emptyset \\ \bar{\bar{K}}_2 - \bar{\bar{K}}_1 & \text{if } K \cap [0, 1] = \emptyset \text{ and } K \cap [1, 2] \neq \emptyset \end{cases},$$

while the prior coverage function is given by $p_{\Gamma_I}(\gamma) = P(\gamma \in [\theta_1, \theta_2]) = \gamma \mathbb{1}\{\gamma \in [0, 1]\} + (2 - \gamma) \mathbb{1}\{\gamma \in [1, 2]\}$, where P is the distribution with respect to the prior of θ .

With this intuition in mind, let us move to the nonparametric Bayesian approach. This approach is based on a Dirichlet process prior and requires to write $\theta = (\theta_1, \theta_2)'$ as $\theta_i = \phi_i(F)$, for a measurable functional $\phi_i, i = 1, 2$. If we observe realizations of a random vector Y from a distribution F then the Bayesian model writes $Y|F \sim F, F \sim \text{Dir}(n_0, F_0)$. The probability measure F_0 should be chosen such that $\phi_2(F) > \phi_1(F)$, μ -a.s. By using the stick-breaking representation of the Dirichlet process, see Appendix D in the Supplement, the prior capacity functional of $\Gamma_I = [\phi_1(F), \phi_2(F)]$ is given by

$$T_{\Gamma_I}(K) = P \left\{ K \cap \left[\phi_1 \left(\sum_{j=1}^{\infty} \alpha_j \delta_{\xi_j} \right), \phi_2 \left(\sum_{j=1}^{\infty} \alpha_j \delta_{\xi_j} \right) \right] \neq \emptyset \right\}, \quad K \in \mathcal{K}, \quad (4.4)$$

where $\{\xi_j\}_{j \geq 1}$ are independent draws from F_0 , δ_{ξ_j} denotes the Dirac mass in ξ_j , $\alpha_j = v_j \prod_{l=1}^{j-1} (1 - v_l)$ with $\{v_l\}_{l \geq 1}$ independent draws from a Beta distribution $\mathcal{B}e(1, n_0)$ and $\{v_j\}_{j \geq 1}$ are independent of $\{\xi_j\}_{j \geq 1}$. In Appendix D we recall how to simulate $\phi_i(F)$, $i = 1, 2$, from the prior and posterior distribution by using this representation. The prior coverage function of Γ_I is: for every $\gamma \in \Gamma$,

$$p_{\Gamma_I}(\gamma) = P(\gamma \in [\phi_1(F), \phi_2(F)]) = P\left(\phi_1\left(\sum_{j=1}^{\infty} \alpha_j \delta_{\xi_j}\right) \leq \gamma \leq \phi_2\left(\sum_{j=1}^{\infty} \alpha_j \delta_{\xi_j}\right)\right) \quad (4.5)$$

where we have taken $K = \{\gamma\}$ and P is the prior distribution of θ , $\{\alpha_j\}$ and $\{\xi_j\}$. In general we do not have an analytic form for T_{Γ_I} and p_{Γ_I} but we have a perfect knowledge of them since we can easily simulate from T_{Γ_I} and p_{Γ_I} by using the stick-breaking representation of the Dirichlet process.

After observing an n -sample of Y , (y_1, \dots, y_n) , one computes the posterior distribution of F as

$$F|y_1, \dots, y_n \sim \mathcal{D}\left(n_0 + n, \frac{n_0}{n_0 + n}F_0 + \frac{n}{n_0 + n}F_n\right),$$

where $F_n(\cdot) := \frac{1}{n} \sum_j \delta_{y_j}(\cdot)$ denotes the empirical cumulative distribution. If the true data distribution F is such that $\phi_2(F) > \phi_1(F)$ then the same is true for the distribution F generated by the posterior.

The posterior capacity functional is denoted by $T_{\Gamma_I}(K|\{y_i\}_{i=1}^n)$ and given by: $\forall K \in \mathcal{K}$,

$$T_{\Gamma_I}(K|\{y_i\}_{i=1}^n) = P \left\{ K \cap \left[\phi_1\left(\rho \sum_{j=1}^n \beta_j \delta_{y_j} + (1-\rho) \sum_{j=1}^{\infty} \alpha_j \delta_{\xi_j}\right), \phi_2\left(\rho \sum_{j=1}^n \beta_j \delta_{y_j} + (1-\rho) \sum_{j=1}^{\infty} \alpha_j \delta_{\xi_j}\right) \right] \neq \emptyset \right\},$$

where $\{\alpha_j\}$, $\{\delta_{\xi_j}\}$, $\{\delta_{y_j}\}$ and $\{\xi_j\}$, are as above, ρ is drawn from a Beta distribution $\mathcal{B}e(n, n_0)$ independently of the other quantities and $(\beta_1, \dots, \beta_n)$ are drawn from a Dirichlet distribution with parameters $(1, \dots, 1)$ on the simplex S_{n-1} of dimension $(n-1)$. For every $\gamma \in \Gamma$, the posterior coverage function $p_{\Gamma_I}(\gamma|\{y_i\}_{i=1}^n) = P(\gamma \in [\phi_1(F), \phi_2(F)]|\{y_i\}_{i=1}^n)$ is: $\forall \gamma \in \Gamma$,

$$P \left(\phi_1\left(\rho \sum_{j=1}^n \beta_j \delta_{y_j} + (1-\rho) \sum_{j=1}^{\infty} \alpha_j \delta_{\xi_j}\right) \leq \gamma \leq \phi_2\left(\rho \sum_{j=1}^n \beta_j \delta_{y_j} + (1-\rho) \sum_{j=1}^{\infty} \alpha_j \delta_{\xi_j}\right) \right). \quad (4.6)$$

For simplicity we have presented only the case where the condition $A(\theta, \gamma) \in A_0$ writes as $\gamma \in [\phi_1(F), \phi_2(F)]$ but our nonparametric method can be generalized to the case where Γ_I is not an interval. In that case, if a Dirichlet process prior is specified for F , the prior capacity functional of Γ_I is given by: $\forall K \in \mathcal{K}$,

$$T_{\Gamma_I}(K) = P(K \cap \Gamma_I \neq \emptyset) = P \left(K \cap \left\{ \gamma \in \Gamma; A\left(\phi\left(\sum_{j=1}^{\infty} \alpha_j \delta_{\xi_j}\right), \gamma\right) \in A_0 \right\} \neq \emptyset \right) \quad (4.7)$$

and the posterior capacity functional is: $\forall K \in \mathcal{K}$,

$$T_{\Gamma_I}(K|\{y_i\}_{i=1}^n) = P(K \cap \Gamma_I \neq \emptyset|\{y_i\}_{i=1}^n) =$$

$$P \left(K \cap \left\{ \gamma \in \Gamma; A(\phi(\rho \sum_{j=1}^n \beta_j \delta_{y_j} + (1 - \rho) \sum_{j=1}^{\infty} \alpha_j \delta_{\xi_j}), \gamma) \in A_0 \right\} \neq \emptyset \middle| \{y_i\}_{i=1}^n \right), \quad (4.8)$$

where $\{\alpha_j\}$, $\{\delta_{\xi_j}\}$, $\{\delta_{y_j}\}$, $\{\xi_j\}$, ρ and $(\beta_1, \dots, \beta_n)$ are as above.

Once the posterior capacity functional is available, an estimator for $\Gamma_I(\theta)$ can be easily constructed. One possibility is to fix θ equal to its posterior mean or median, denote it by $\hat{\theta}$, and take the corresponding $\Gamma_I(\hat{\theta})$ as an estimator for Γ_I . In alternative, one could construct a closed set C_α that satisfies the following condition

$$P(\Gamma_I(\theta) \subset C_\alpha | \{y_i\}_{i=1}^n) \geq \alpha \quad (4.9)$$

for some $\alpha \in [0, 1]$. This is the usual posterior credible region. A similar estimator is proposed *e.g.* by Liao and Simoni (2019) and Chen et al. (2018). We remark that the probability in (4.9) is determined by the posterior Dirichlet process and is the *posterior containment functional* evaluated at C_α .

4.2 Examples

In this section we provide two examples where we use our proposed Bayesian nonparametric method described in section 4.1. Other two examples are developed in Appendix B in the Supplement. For simplicity, we only focus on the prior and posterior coverage functions, which are easy to represent graphically.

Example 2 (Interval Censored Data.) *This example is motivated by interval responses in survey data. Let Y be the real random variable of interest that is unobserved but is known to lie in the interval $[Y_1, Y_2]$ a.s. with respect to the sampling distribution, where Y_1 and Y_2 are two observable real random variables. The probability distributions of Y_1 and Y_2 are unknown and denoted by F_1 and F_2 , respectively. We denote with $\mathbf{E}_{F_1}(Y_1)$ and $\mathbf{E}_{F_2}(Y_2)$ the expectation taken with respect to F_1 and F_2 , respectively. Let $\gamma := \mathbf{E}(Y) \in \Gamma = \mathbb{R}$ be the parameter of interest. Since $Y \in [Y_1, Y_2]$ a.s., the condition $A(\theta, \gamma) \in A_0$ takes the form*

$$\mathbf{E}_{F_1}(Y_1) \leq \gamma \leq \mathbf{E}_{F_2}(Y_2) \quad (4.10)$$

a.s. with respect to the prior distribution μ of θ , where $\theta := \phi(F) = (\mathbf{E}_{F_1}(Y_1), \mathbf{E}_{F_2}(Y_2))'$ and $F := (F_1, F_2)'$ is the joint distribution of (Y_1, Y_2) . More precisely, for every $\gamma \in \Gamma$, $A(\theta, \gamma) = [\mathbf{E}_{F_1}(Y_1) - \gamma, \gamma - \mathbf{E}_{F_2}(Y_2)]$ and $A_0 = (-\infty, 0] \times (-\infty, 0]$ which is an element of the Borel σ -field of subsets of $\Phi = \mathbb{R}^2$. Hence, $\Gamma_I = [\mathbf{E}_{F_1}(Y_1), \mathbf{E}_{F_2}(Y_2)]$ is a random closed set, the parameter γ is partially identified and our object of interest becomes the identified set Γ_I .

We compute the prior and posterior coverage function of the identified set Γ_I by specifying a Dirichlet process prior for θ . Let us assume that $Y_1 \perp Y_2 | F$. For F_0^1 and F_0^2 two probability measures, the Bayesian hierarchical model is

$$\begin{aligned} (y_{11}, \dots, y_{1n}) | F &\sim iid \ F_1 \\ (y_{21}, \dots, y_{2n}) | F &\sim iid \ F_2 \\ F_1 &\sim Dir(n_0^1, F_0^1), \quad n_0^1 \in \mathbb{R}_+ \\ F_2 &\sim Dir(n_0^2, F_0^2), \quad n_0^2 \in \mathbb{R}_+, \end{aligned} \quad (4.11)$$

where (y_{11}, \dots, y_{1n}) and (y_{21}, \dots, y_{2n}) denote two n -samples of realisations of Y_1 and Y_2 , respectively. If the probability measures F_0^1 and F_0^2 have disjoint supports, that is, $\max \text{Supp}(F_0^1) < \min \text{Supp}(F_0^2)$ then, $\mathbf{E}_{F_2}(Y_2) > \mathbf{E}_{F_1}(Y_1)$ μ -a.s. For every $\gamma \in \Gamma$, let $g_\theta(\gamma) := \mathbb{1}\{\gamma \in [\mathbf{E}_{F_1}(Y_1), \mathbf{E}_{F_2}(Y_2)]\}$. For every $\gamma \in \Gamma$ the prior coverage function is $\mathbf{E}(g_\theta(\gamma))$, where the expectation is taken with respect to the prior μ of F , and can be represented as, $\forall \gamma \in \Gamma$,

$$p_{\Gamma_I}(\gamma) = P\left(\gamma \in [\mathbf{E}_{F_1}(Y_1), \mathbf{E}_{F_2}(Y_2)]\right) = P\left(\sum_j \alpha_j^1 \xi_j^1 \leq \gamma \leq \sum_j \alpha_j^2 \xi_j^2\right)$$

where, for $i = 1, 2$, $\{\xi_j^i\}_{j \geq 1}$ are independent draws from F_0^i , $\alpha_j^i = v_j^i \prod_{l=1}^{j-1} (1 - v_l^i)$ with $\{v_l^i\}_{l \geq 1}$ independent draws from a Beta distribution $\mathcal{B}e(1, n_0^i)$ and $\{v_j^i\}_{j \geq 1}$ are independent of $\{\xi_j^i\}_{j \geq 1}$. For $i = 1, 2$, let $\mathcal{F}_i(\cdot)$ denote the cumulative distribution function of $\mathbf{E}_{F_i}(Y_i)$. Hence, $p_{\Gamma_I}(\gamma)$ takes the following values: (i) $\forall \gamma \notin [\min \text{Supp}(F_0^1), \max \text{Supp}(F_0^2)]$, $p_{\Gamma_I}(\gamma) = 0$; (ii) $\forall \gamma \in [\max \text{Supp}(F_0^1), \min \text{Supp}(F_0^2)]$, $p_{\Gamma_I}(\gamma) = 1$; (iii) $\forall \gamma \in \text{Supp}(F_0^1)$, $p_{\Gamma_I}(\gamma) = P\left(\sum_j \alpha_j^1 \xi_j^1 \leq \gamma\right) \equiv \mathcal{F}_1(\gamma)$; (iv) $\forall \gamma \in \text{Supp}(F_0^2)$, $p_{\Gamma_I}(\gamma) = P\left(\gamma \leq \sum_j \alpha_j^2 \xi_j^2\right) \equiv 1 - \mathcal{F}_2(\gamma)$.

The posterior distributions of F_1 and F_2 are $F_i | (y_{i1}, \dots, y_{in}) \sim Dir(n_*^i, F_*^i)$, $i = 1, 2$, where for

$i = 1, 2$, $n_*^i = n_0^i + n$, $F_*^i = \frac{n_0^i}{n_0^i + n} F_0^i + \frac{n}{n_0^i + n} F_n^i$ and $F_n^i(\cdot) := \frac{1}{n} \sum_j \delta_{y_{ij}}(\cdot)$ is the empirical distribution of the sample (y_{i1}, \dots, y_{in}) . The posterior coverage function of Γ_I is, $\forall \gamma \in \Gamma$

$$p_{\Gamma_I}(\gamma | \{y_{1i}\}, \{y_{2i}\}) = P\left(\gamma \in [\mathbf{E}_{F_1}(Y_1), \mathbf{E}_{F_2}(Y_2)] \mid \{y_{1i}\}, \{y_{2i}\}\right)$$

and can be represented as:

$$P\left(\rho_1 \sum_{j=1}^n \beta_{1j} y_{j1} + (1 - \rho_1) \sum_{j=1}^{\infty} \alpha_j^1 \xi_j^1 \leq \gamma \leq \rho_2 \sum_{j=1}^n \beta_{2j} y_{j2} + (1 - \rho_2) \sum_{j=1}^{\infty} \alpha_j^2 \xi_j^2\right),$$

where, for $i = 1, 2$, α_j^i and ξ_j^i are as above, ρ_i is drawn from a Beta distribution $\mathcal{B}e(n, n_0^i)$ independently of the other quantities and $(\beta_{i1}, \dots, \beta_{in})$ are drawn from a Dirichlet distribution with parameters $(1, \dots, 1)$ on the simplex S_{n-1} of dimension $(n - 1)$.

A simulation exercise allows to visualize the prior and posterior coverage functions of Γ_I . We generate an n -sample of realizations of Y_1 and Y_2 from the following distributions¹:

$$Y_1 \sim \mathcal{N}(0, 0.1), \quad Y_2 \sim \mathcal{N}(5, 0.1).$$

The parameters are fixed as follows: $n = 1000$, $n_0^1 = 10$, $n_0^2 = 20$, $F_0^1 = \mathcal{N}(0, 1)$ and $F_0^2 = \mathcal{N}(10, 1)$. The supports of F_0^1 and F_0^2 are not disjoint. However, since the tails of a normal density function are very thin, the prior probability that $\mathbf{E}_{F_2}(Y_2) < \mathbf{E}_{F_1}(Y_1)$ is very small. The true identified set in our simulation is $[\mathbf{E}_{F_1}(Y_1), \mathbf{E}_{F_2}(Y_2)] = [0, 5]$.

In Figure 1 we represent the subset $[-3, 12]$ of Γ on the horizontal axis and we evaluate $p_{\Gamma_I}(\cdot)$ and $p_{\Gamma_I}(\cdot | \{y_{1i}\}, \{y_{2i}\})$ over a grid of values on this intervals. Then, we draw 1000 times from the prior and posterior distribution of (F_1, F_2) and for every value γ in the grid of $[-3, 12]$ we count the number of simulated Γ_I s that contain this value γ . Figure 1 shows the prior and posterior coverage functions for each value of $\gamma \in [-3, 12]$. Figure 2 displays the intervals Γ_I drawn from the prior and posterior distributions (on the vertical axis) against the true interval $[0, 5]$ (on the horizontal axis). It also plots the true interval $[0, 5]$ against itself in red. Both figures show that the posterior concentrates on the true $\Gamma_I = [0, 5]$.

¹This data generating process is the same used in Liao and Jiang (2010) in their Example 5.1.

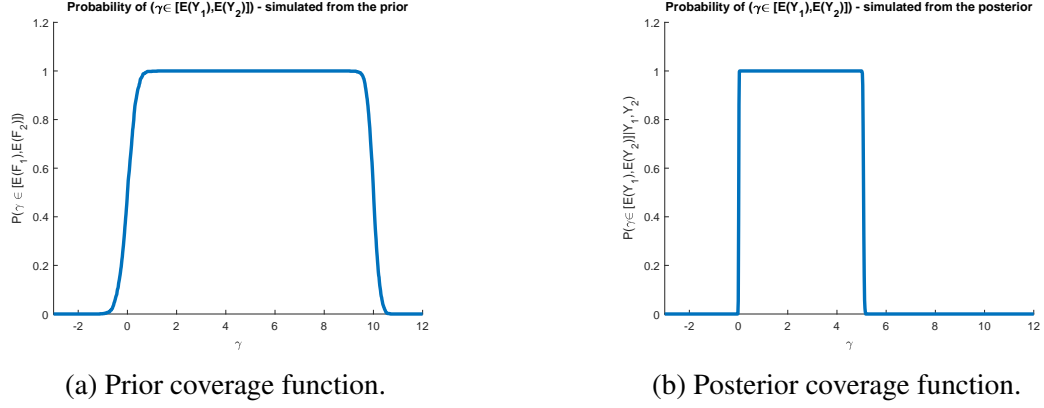


Figure 1: Interval Censored Data. Prior and posterior coverage functions $p_{\Gamma_I}(\cdot)$ and $p_{\Gamma_I}(\cdot | \{y_{1i}\}, \{y_{2i}\})$ of Γ_I . The true Γ_I is $[0, 5]$.




Figure 2: Interval Censored Data. Representation of the intervals $\Gamma_I(\theta) = [\mathbf{E}_{F_1}(Y_1), \mathbf{E}_{F_2}(Y_2)]$ drawn from the prior (panel 2a) and posterior (panel 2b) against the true interval $\Gamma_I = [0, 5]$ (in red).

Example 3 (Linear Regression with errors in Regressors.) *This is the well-known linear errors-in-variables structural model considered in Frisch (1934) and Klepper and Leamer (1984). For simplicity, we focus here on the univariate linear regression model. Let Y be an observable random variable satisfying the model $Y = \gamma\xi + u$ where ξ is an unobservable random variable such that $\mathbf{E}(\xi) = 0$, $\text{Var}(\xi) = \tau^2$ and for which only realizations affected by an error are available: $Z = \xi + v$. The error terms (u, v) are zero-mean jointly distributed random variables, independent of ξ , and with a variance-covariance matrix Σ which may be diagonal.*

Suppose $\Sigma = \text{diag}(\sigma_u^2, \sigma_v^2)$ for simplicity and that $\mathbf{E}(\xi)$ is known, so that the structural parameters are $\gamma, \sigma_w^2, \sigma_v^2$ and τ^2 while ξ is the incidental parameter. Denote $\sigma_{yz} := \text{Cov}(Y, Z)$, $\sigma_{zz} := \text{Var}(Z)$ and $\sigma_{yy} := \text{Var}(Y)$. Due to endogeneity of Z , the parameter γ lies in the identified set $\Gamma_I = [\theta_1, \theta_2]$

given by

$$\begin{aligned} [\theta_1, \theta_2] &= \left[\min \left(\frac{\sigma_{yz}}{\sigma_{zz}}, \frac{\sigma_{yy}}{\sigma_{yz}} \right), \max \left(\frac{\sigma_{yz}}{\sigma_{zz}}, \frac{\sigma_{yy}}{\sigma_{yz}} \right) \right] \\ &= \left[\min \left(\frac{\gamma\tau^2}{\tau^2 + \sigma_v^2}, \gamma + \frac{\sigma_u^2}{\gamma\tau^2} \right), \max \left(\frac{\gamma\tau^2}{\tau^2 + \sigma_v^2}, \gamma + \frac{\sigma_u^2}{\gamma\tau^2} \right) \right], \end{aligned} \quad (4.12)$$

where $\frac{\sigma_{yz}}{\sigma_{zz}}$ is the coefficient of the regression line of Y on Z and $\frac{\sigma_{yz}}{\sigma_{yy}}$ is the coefficient of the reverse regression line of Z on Y , both without intercept. This can be seen by running the two regressions $y = \gamma(Z - v) + u$ and $Z = \frac{1}{\gamma}(Y - u) + v$. The first regression gives: $\gamma = \sigma_{yz}/(\sigma_{zz} - \sigma_v^2) \geq \sigma_{yz}/\sigma_{zz}$ while the second regression gives: $\gamma = (\sigma_{yy} - \sigma_u^2)/\sigma_{yz} \leq \sigma_{yy}/\sigma_{yz}$.

We specify a Dirichlet process prior on the joint probability distribution F_{yz} of (Y, Z) . In the simulation exercise we generate the data as: for $i = 1, \dots, n$,

$$\xi_i \sim i.i.d. \mathcal{N}(0, 1), \quad (u_i, v_i)' \sim i.i.d. \mathcal{N}_2(0, I_2), \quad (4.13)$$

$$Z_i = \xi_i + v_i,$$

$$Y_i = \gamma\xi_i + u_i, \quad \gamma = 1,$$

where I_2 is the 2-dimensional identity matrix. Since $\gamma > 0$ the identified set is $[\frac{\sigma_{yz}}{\sigma_{zz}}, \frac{\sigma_{yy}}{\sigma_{yz}}] = [1/2, 2]$.

We specify the prior on F_{yz} as $F_{yz} \sim \text{Dir}(n_0, F_0)$ with $n_0 = 20$ and

$$F_0 = \mathcal{N}_2 \left(\begin{pmatrix} 0 \\ 0 \end{pmatrix}, \begin{pmatrix} 2 & 0.9 \\ 0.9 & 2 \end{pmatrix} \right). \quad (4.14)$$

The sample size is $n = 1000$ and we draw 1000 intervals Γ_I from the prior and posterior distribution of F . The results are shown in Figures 3 and 4.

5 Marginal Identification

In sections 3 and 4 we have considered statistical models where no parameter is marginalized out. We refer to these models as full models. When the parameter of interest is a sub-parameter γ of the whole

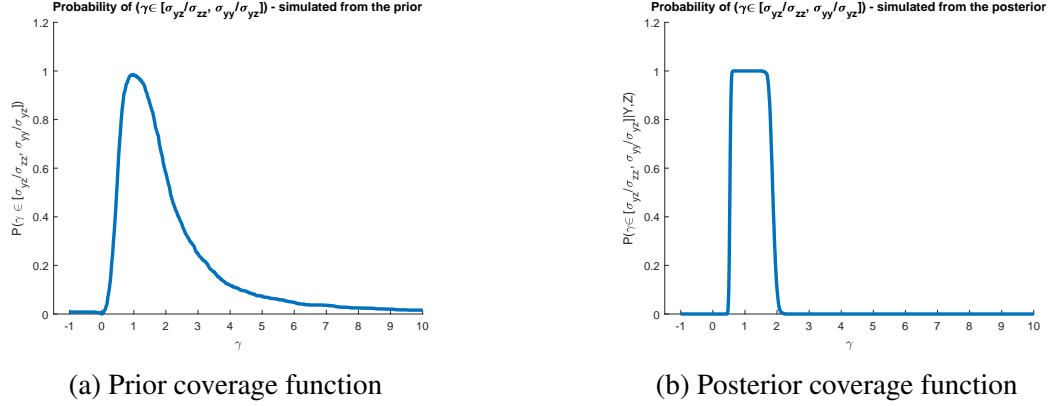


Figure 3: Linear Regression with errors in Regressors. Prior and posterior coverage functions $p_{\Gamma_I}(\cdot)$ and $p_{\Gamma_I}(\cdot|\{y_i\}, \{z_i\})$ of Γ_I . The true Γ_I is $[1/2, 2]$.




Figure 4: Linear Regression with errors in Regressors. Representation of the intervals $\Gamma_I(\theta) = \left[\frac{\sigma_{yz}}{\sigma_{zz}}, \frac{\sigma_{yy}}{\sigma_{yz}} \right]$ drawn from the prior (panel 4a) and posterior (panel 4b) against the true interval $[1/2, 2]$ (in red).

model parameter one might want to perform the analysis by getting rid of the parameters of the model that are not of interest. Therefore, it is natural to examine the marginal model in this sub-parameter.

5.1 Marginal Identification of nonidentified sub-parameters

Let $\theta := (\beta, \gamma)$ denote the whole parameter of the model where β is identified and γ is the parameter of interest that is nonidentified. For instance, γ is the parameter related to a latent variable as in section 3.1. One can specify the prior for θ as $\mu(\theta) = \mu(\beta|\gamma)\mu(\gamma)$. Hence, the marginal model is obtained by

integrating out the parameter β in the original model with respect to the prior $\mu(\beta|\gamma)$:

$$P^\gamma(E) = \int P^\beta(E)\mu(d\beta|\gamma), \quad \forall E \in \mathcal{X},$$

where $P^\gamma(\cdot)$ denotes the integrated sampling distribution which depends on γ . The corresponding Lebesgue density function (or marginal likelihood) writes $p(x|\gamma) = \int p(x|\beta, \gamma)\mu(\beta|\gamma)d\beta$. The predictive density $p(x)$, obtained by integrating out γ from $p(x|\gamma)$ with respect to the prior, is the same as in the full model and the marginal posterior $\mu(\gamma|x)$ of γ is obtained by marginalizing the joint posterior $\mu(\beta, \gamma|x)$ with respect to β .

Example 4 (Classical model of hyperparameter) *Let us consider the following Bayesian model $x_i|\beta, \gamma \sim \text{iid } \mathcal{N}(\beta, \sigma^2)$, $i = 1, \dots, n$, $\beta|\gamma \sim \mathcal{N}(\gamma, \sigma_0^2)$, $\gamma \sim \mathcal{N}(\gamma_0, \tau_0^2)$, where γ is the hyperparameter of the prior distribution and $\sigma^2, \sigma_0^2, \gamma_0$, and τ_0^2 are known parameters. The parameter γ is unidentified in the sampling model. The minimal sufficient σ -field $\mathcal{A}\mathcal{X}$ is almost surely equal to the σ -field generated by β . The marginal model is $(x_1, \dots, x_n)'|\gamma \sim \mathcal{N}(\gamma\iota, \sigma^2 I_n + \sigma_0^2 \iota \iota')$, where $\iota = (1, \dots, 1)'$ is $n \times 1$ and I_n is the n -dimensional identity matrix. Therefore, γ is identified in the marginal model.*

Note that, in the previous example, even if γ is identified in the marginal model it is not exactly estimable. In fact, Theorem 2.3 does not apply because the marginal model is not *i.i.d.* Exact estimability would hold only if the conditional distribution of γ given β was a degenerated Dirac measure on a deterministic function of β .

In the setting of Example 4, β can be interpreted as an heterogeneity parameter whose distribution depends on an unidentified parameter γ , see *e.g.* Heckman and Singer (1984). In the frequentist setting, the conditional distribution of $\beta|\gamma$ is part of the data generating process while in the Bayesian setting the conditional distribution of $\beta|\gamma$ is the prior.

The next theorem considers the asymptotic behavior of a sub-parameter which might be nonidentified in the full model. It states that the posterior mean of a sub-parameter c converges Π -*a.s.* to the conditional prior mean given the identified parameter. Remark that the *a.s.* in the theorem is with respect to the joint distribution.

Theorem 5.1 *Let us consider a Bayesian model $\{\Theta \times X, \mathcal{A} \otimes \mathcal{X}, \Pi\}$ with a filtration $\mathcal{X}_n \rightarrow \mathcal{X}_\infty$ and*

where the identified parameter \mathcal{AX}_∞ is asymptotically exactly estimable, that is, $\mathcal{AX}_\infty \subset \overline{\mathcal{X}_\infty}$. Let c be an integrable function defined on \mathcal{A} , then: $\mathbf{E}(c|\mathcal{X}_n) \rightarrow \mathbf{E}(c|\mathcal{AX}_\infty)$, $\Pi - a.s.$

5.2 Marginal identification in partially identified models

Consider now the partially identified model of section 4. Let Y be an observable random variable with distribution F . The parameter F is identified and suppose that there is another parameter θ of the model that is identified and that can be written as $\theta = \phi(F)$ for some functional ϕ . The parameter of interest is denoted by γ and is related to θ by relation (4.1): $A(\theta, \gamma) \in A_0 \subset \Phi$. Let Γ be provided with a σ -field \mathcal{G} . Hence, the parameter space is $(\Theta \times \Gamma, \mathcal{A} \otimes \mathcal{G})$. By using the structural relation (4.1) we now specify a restricted prior $\mu(\gamma|\theta)$ for γ conditional on θ . In particular, $\mu(\gamma|\theta)$ has support equal to the set of the γ s that satisfy the constraint $A(\theta, \gamma) \in A_0$ in (4.1) for a given $\theta \in \Theta$.² The marginal posterior of γ is: $\forall \Gamma_1 \in \mathcal{G}$,

$$\mu(\Gamma_1|y) = \int_{\Gamma_1} \int_{\Theta} \mu(\theta|y) \mu(\gamma|\theta) d\theta d\gamma = \int_{\Gamma_1} d\gamma \int_{\{F; \phi(F) \in \Theta\}} \mu(\gamma|\phi(F)) \mu(dF|y),$$

where $\mu(\theta|y)$ denotes the posterior distribution of θ and $\mu(dF|y)$ denotes the posterior distribution of F . In the second equality we have written the integral in terms of $\mu(dF|y)$ to stress the fact that the prior distribution of θ is recovered from the Dirichlet process prior for F as described in section 4. It is clear that γ is identified in the marginal model since its marginal prior distribution is updated by the data. In addition, Theorem 5.1 above applies also to the case of partially identified models.

While the marginal posterior density $\mu(\gamma|y)$ of γ is not usually known in closed-form – in particular if θ is infinite dimensional – one can easily simulate from it. For this, one first simulates θ given y from $\mu(\theta|y)$ and then, for each draw of θ , one simulates γ given θ from $\mu(\gamma|\theta)$. This simulation scheme produces draws from $\mu(\gamma|y)$ and this is due to the lack of identification of γ which implies that $\gamma \perp y|\theta$, see section 2.3. Having a marginal posterior distribution $\mu(\gamma|y)$ of the parameter of interest γ is important for instance in a decision problem setting. In fact, knowledge of $\mu(\gamma|y)$ allows to select the most likely value of γ or the region inside Γ_I with the highest posterior probability. Such a selection is

²In alternative, we may relax this constraint on the support of $\mu(\gamma|\theta)$ into a constraint on the hyperparameter of the distribution of γ , as illustrated by the examples below.

clearly affected by the choice of the prior on γ .

5.3 Examples

In this section we develop further Example 3 of section 4.2 by endowing the parameter γ with a conditional prior distribution given $\theta = \phi(F)$ that we denote by $\mu_\gamma^F: \gamma|\theta \sim \mu_\gamma^F := \mu(\gamma|\phi(F))$. Additional examples are developed in Appendix C in the Supplement. In our simulations we consider four different specifications for μ_γ^F , where the hyperparameters a_0 and b_0 are specified in each specific example:

(I). $\gamma|F \sim \mathcal{N}(\gamma_0, \tau_0^2)$, $\gamma_0 = \frac{\tilde{\gamma}_0^1 + \tilde{\gamma}_0^2}{2c_0}$, $\tau_0 \in \mathbb{R}$ and we discard the draws of γ that do not belong to the interval $[a_0, b_0]$;

(II). $\gamma|F \sim \mathcal{N}(0, \sigma_0^2)$ truncated to the interval $[a_0, b_0]$;

(III). $\gamma|F \sim \mathcal{U}[a_0, b_0]$ (flat prior);

(IV). $\gamma|F \sim \mathcal{Be}(a_0, b_0, p, q)$, that is, a Beta prior distribution with support $[a_0, b_0]$ and shape parameters p and q . The corresponding probability density function is:

$$\mu(\gamma|F) = \frac{(\gamma - a_0)^{p-1} (b_0 - \gamma)^{q-1}}{B(p, q) (b_0 - a_0)^{p+q-1}}$$

where $B(p, q)$ is the beta function.

Example 3 (Linear Regression with errors in Regressors (*continued*)). Suppose that we are not only interested in the identified region $\Gamma_I := [\theta_1, \theta_2]$, where θ_1, θ_2 are defined in 4.12, but also in the parameter γ itself. The marginal posterior distribution of γ is informative about the areas of the identified region Γ_I where the parameter is more likely. The prior distribution of $\theta := (\theta_1, \theta_2)$ is obtained from a Dirichlet process prior for the joint probability distribution F_{yz} of (Y, Z) denoted by μ_F , see section 4. The Bayesian hierarchical model is as in (4.13)-(4.14) completed with the specification of a prior for γ conditional on F_{yz} : $\gamma|F_{yz} \sim \mu_\gamma^{F_{yz}}$. In our simulation exercise we consider the four specifications (I)-(IV) for $\mu_\gamma^{F_{yz}}$ given above with: $a_0 = \min\left(\frac{\sigma_{yz}}{\sigma_{zz}}, \frac{\sigma_{yy}}{\sigma_{yz}}\right)$, $b_0 = \max\left(\frac{\sigma_{yz}}{\sigma_{zz}}, \frac{\sigma_{yy}}{\sigma_{yz}}\right)$, $c_0 = 1$, and $\tilde{\gamma}_0^1 = a_0$, $\tilde{\gamma}_0^2 = b_0$. Since γ is nonidentified in the full model, but identified in the marginal model, its posterior

distribution depends on the data only through F . This means that the moments σ_{yz} , σ_{zz} and σ_{yy} in the prior for γ , computed from the Dirichlet process prior for F_{yz} are replaced with the posterior means of σ_{yz} , σ_{zz} and σ_{yy} in the posterior distribution, computed from the posterior of the Dirichlet process for F_{yz} .

We generate an n -sample of observations of (Y_1, Y_2) as in (4.13). The parameters are fixed as follows: $n = 1000$, F_0 is specified as in (4.14), $\tau_0^2 = 1$, $\sigma_0^2 = 2$, $p = 1$ and $q = 0.5$. The true identified set is $\Gamma_I = [1/2, 2]$.


We draw 1000 times from the marginal prior and posterior distributions of γ . The simulation scheme is the following: for each $1 \leq j \leq 1000$, draw $F_{yz}^{(j)}$ from the prior $\mu(F_{yz})$ (resp. the posterior $\mu(F_{yz}|\{y_i, z_i\}_{i=1}^n)$), compute $\theta_i^{(j)}$, $i = 1, 2$ and draw $\gamma^{(j)}$ from $\mu_{\gamma}^{F_{yz}^{(j)}}$ (resp. $\mu_{\gamma}^{F_{yz}^{(j)}}(\gamma|F_{yz}^{(j)}, \{y_i, z_i\}_{i=1}^n)$). Figure 5 shows the histograms of the marginal prior (in blue) and posterior (in red) distribution of γ . Each panel corresponds to one of the four specifications for $\mu_{\gamma}^{F_{yz}}$. We see that the marginal posterior distribution is much more concentrated on the true identified set than the corresponding prior.

6 Conclusions


This paper studies theoretical properties and implementation of the Bayesian approach in various models that lack identification. As examples of unidentified models, we analyse nonparametric models with heterogeneity modeled either as a Gaussian process or as a Dirichlet process where the parameter of interest is the (hyper)parameter of the heterogeneity distribution which is unidentified. We also analyse unidentified latent variable and partially identified models.

In partially identified models we propose to construct the prior and posterior of the identified set through the prior and posterior capacity functionals. The prior capacity functional is obtained as the transformation of a Dirichlet process, so that our approach is completely nonparametric. The proposed procedure is appealing since, even if the posterior capacity functional has a complicated expression, simulating from it is simple.


Finally, we discuss models that have some parameters that are identified and others that are not. For these models, we show that the parameter that is unidentified or partially identified in the full model is identified in the marginal model.




(a) $\gamma|F_{yz} \sim \mathcal{N}(\gamma_0, 1)$ by discarding the draws that are not in $[\theta_1, \theta_2]$.



(b) $\gamma|F_{yz} \sim \mathcal{N}(0, 2)$ truncated to $[\theta_1, \theta_2]$.



(c) $\gamma|F_{yz} \sim \mathcal{U}[\mathbf{E}_{F_1}(Y_1), \mathbf{E}_{F_2}(Y_2)]$.



(d) $\gamma|F_{yz} \sim \text{Beta}(\mathbf{E}_{F_1}(Y_1), \mathbf{E}_{F_2}(Y_2), 2, 2)$.

Figure 5: Linear Regression with errors in Regressors. Histograms of the prior (in blue) and posterior (in red) probability distributions. The true identified set is $\Gamma_I = [1/2, 2]$.

References

- Berger, J. (1985), *Statistical decision theory and Bayesian analysis*, Springer-Verlag New York.
- Bontemps, C. and Magnac, T. (2017), ‘Set identification, moment restrictions, and inference’, *Annual Review of Economics* **9**(1), 103–129.
- Chen, X., Christensen, T. and Tamer, E. T. (2018), ‘Monte Carlo confidence sets for identified

- sets', *Econometrica* **86**(6), 1965–2018.
- Chernozhukov, V. and Hong, H. (2003), 'An mcmc approach to classical estimation', *Journal of Econometrics* **115**, 293–346.
- Dellacherie, C. and Meyer, P. (1975), *Probabilité et potentiel*, Paris: Hermann.
- Diaconis, P. and Freedman, D. (1986), 'On the consistency of bayes estimates', *Ann. Statist.* **14**(1), 1–26.
- Drèze, J. H. and Richard, J.-F. (1983), Chapter 9 Bayesian analysis of simultaneous equation systems, Vol. 1 of *Handbook of Econometrics*, Elsevier, pp. 517 – 598.
- Ferguson, T. S. (1973), 'A Bayesian analysis of some nonparametric problems', *Ann. Statist.* **1**(2), 209–230.
- Florens, J. (2002), Inférence Bayésienne non paramétrique et bootstrap., in J. F. . G. S. J.J. Dreesbeke, ed., 'Méthodes Bayésiennes en statistique', Editions TECHNIP, pp. 295–313.
- Florens, J. and Mouchart, M. (1986), 'Exhaustivité, ancillante et identification en statistique Bayésienne', *Annales d'Economie et Statistiques* **4**, 63–93.
- Florens, J., Mouchart, M. and Rolin, J. (1985), 'On two definitions of identification', *Statistics* **16**, 213–218.
- Florens, J., Mouchart, M. and Rolin, J. (1990), *Elements of Bayesian statistics.*, Dekker - New York.
- Florens, J.-P. and Simoni, A. (2012), 'Regularized posteriors in linear ill-posed inverse problems', *Scandinavian Journal of Statistics* **39**(2), 214–235.
- Frisch, R. (1934), 'Statistical confluence analysis by means of complete regression systems', *University Institute of Economics, Oslo* pp. 5–8.
- Giacomini, R. and Kitagawa, T. (2021), 'Robust bayesian inference for set-identified models', *Econometrica* **89**(4), 1519–1556.
- Gustafson, P. (2005), 'On model expansion, model contraction, identifiability and prior information: Two illustrative scenarios involving mismeasured variables', *Statistical Science* **20**, 111 –140.
- Gustafson, P. (2010), 'Bayesian inference for partially identified models', *International Journal of Biostatistics* **17**, 2107–2124.
- Gustafson, P. (2012), 'On the behaviour of Bayesian credible intervals in partially identified models', *Electronic Journal of Statistics* **6**, 2107–2124.
- Heckman, J. and Singer, B. (1984), 'A method for minimizing the impact of distributional assumptions in econometric models for duration data', *Econometrica* **52**(2), 271–320.
- Hoogerheide, L., Kleibergen, F. and van Dijk, H. K. (2007), 'Natural conjugate priors for the instrumental variables regression model applied to the Angrist–Krueger data', *Journal of Econometrics* **138**(1), 63 – 103.
- Husmoller, D. (1994), *Fibre bundles*, Springer-Verlag.

- Kadane, J. (1974), The role of identification in Bayesian theory, in S. Fienberg and A. Zellner, eds, 'Studies in Bayesian Econometrics and Statistics', North Holland.
- Kleibergen, F. and Mavroeidis, S. (2014), 'Identification issues in limited-information bayesian analysis of structural macroeconomic models', *Journal of Applied Econometrics* **29**(7), 1183–1209.
- Kleibergen, F. and van Dijk, H. K. (1998), 'Bayesian simultaneous equations analysis using reduced rank structures', *Econometric Theory* **14**(6), 701–743.
- Klepper, S. and Leamer, E. E. (1984), 'Consistent sets of estimates for regressions with errors in all variables', *Econometrica* **52**(1), 163–183.
- Kline, B. and Tamer, E. (2016), 'Bayesian inference in a class of partially identified models', *Quantitative Economics* **7**(2), 329–366.
- Kolmogorov, A. and Fomin, S. (1975), *Introductory real analysis*, Dover.
- Liao, Y. and Jiang, W. (2010), 'Bayesian analysis in moment inequality models', *Annals of Statistics* **38**, 275–316.
- Liao, Y. and Simoni, A. (2019), 'Bayesian inference for partially identified smooth convex models', *Journal of Econometrics* **211**, 338–360.
- Lindley, D. (1971), *Bayesian statistics: a review*, Philadelphia, SIAM.
- Molchanov, I. (2005), *Theory of random sets*, Springer.
- Molinari, F. (2020), Chapter 5 - Microeconometrics with partial identification, in S. N. Durlauf, L. P. Hansen, J. J. Heckman and R. L. Matzkin, eds, 'Handbook of Econometrics, Volume 7A', Vol. 7 of *Handbook of Econometrics*, Elsevier, pp. 355–486.
- Moon, R. and Schorfheide, F. (2012), 'Bayesian and frequentist inference in partially-identified models', *Econometrica* **80**, 755–782.
- Norets, A. and Tang, X. (2013), 'Semiparametric inference in dynamic binary choice models', *The Review of Economic Studies* **81**(3), 1229–1262.
- Picci, G. (1977), 'Some connections between the theory of sufficient statistics and the identifiability problem', *SIAM Journal on Applied Mathematics* **33**(3), 383–398.
- Poirier, D. J. (1998), 'Revising beliefs in nonidentified models', *Econometric Theory* **14**(4), 483–509.
- San Martín, E. and Gonzáles, J. (2010), 'Bayesian identifiability: contributions to an inconclusive debate', *Chilean Journal of Statistics* **1**, 69–91.
- van Dyk, D. A. and Meng, X.-L. (2001), 'The art of data augmentation', *Journal of Computational and Graphical Statistics* **10**(1), 1–50.