
The Prior Influence Function in Variational Bayes

Ryan Giordano rgiordano@berkeley.edu

Tamara Broderick tbroderick@csail.mit.edu

Michael Jordan jordan@cs.berkeley.edu

Abstract

In Bayesian analysis, the posterior follows from the data and a choice of a prior and a likelihood. One hopes that the posterior is robust to reasonable variation in the choice of prior, since this choice is made by the modeler and is often somewhat subjective. A different, equally subjectively plausible choice of prior may result in a substantially different posterior, and so different conclusions drawn from the data. Were this to be the case, our conclusions would not be robust to the choice of prior. To determine whether our model is robust, we must quantify how sensitive our posterior is to perturbations of our prior.

Variational Bayes (VB) methods are fast, approximate methods for posterior inference. As with any Bayesian method, it is useful to evaluate the robustness of a VB approximate posterior to changes in the prior. In this paper, we derive VB versions of classical non-parametric local robustness measures. In particular, we show that the influence function of Gustafson (2000) has a simple, easy-to-calculate closed form expression for VB approximations. We then demonstrate how local robustness measures can be inadequate for non-local prior changes, such as replacing one prior entirely with another. We propose a simple approximate non-local robustness measure and demonstrate its effectiveness on a simulated data set.

1 Local robustness and the influence function

Bayesian robustness studies how changes to the model (i.e., the prior and likelihood) and to the data affect the posterior. If important aspects of the posterior are meaningfully sensitive to subjectively reasonable perturbations of the inputs, then the posterior is “non-robust” to these perturbations. In this paper, we focus on quantifying the sensitivity of posterior means to perturbations of the prior – either infinitesimally mixing or completely replacing the original prior with another “contaminating prior”. Our methods allow fast estimation of sensitivity to any contaminating prior without re-fitting the model. We follow and extend the work of Gustafson (1996) and Gustafson (2000) to variational Bayes and to approximate non-local measures of sensitivity. For a more general review of Bayesian robustness, see Berger et al. (2000).

We will now define some terminology. Denote our N data points by $x = (x_1, \dots, x_N)$ with $x_n \in \mathbb{R}^D$. Denote our parameter by the vector $\theta \in \mathbb{R}^K$. We will suppose that we are interested in the robustness of our prior to a scalar parameter ϵ where our prior can be written as $p(\theta|\epsilon)$. Let p_ϵ^x denote the posterior distribution of θ with prior given by ϵ and conditional on x , as given by Bayes’ Theorem:
$$p_\epsilon^x(\theta) := p(\theta|x, \epsilon) = \frac{p(x|\theta)p(\theta|\epsilon)}{p(x)}.$$

A typical end product of a Bayesian analysis might be a posterior expectation of some function, $\mathbb{E}_{p_\epsilon^x}[g(\theta)]$, which is a functional of $g(\theta)$ and $p_\epsilon^x(\theta)$. Local robustness considers how much $\mathbb{E}_{p_\epsilon^x}[g(\theta)]$ changes locally in response to small perturbations in the value of ϵ (Gustafson, 2000). In the present

work, we consider mixing our original prior, $p_0(\theta)$, with some known alternative functional form, $p_c(\theta)$:

$$p(\theta|\epsilon) = (1 - \epsilon)p_0(\theta) + \epsilon p_c(\theta) \text{ for } \epsilon \in [0, 1]. \quad (1)$$

This is known as epsilon contamination (the subscript c stands for ‘‘contamination’’), and its construction guarantees that the perturbed prior is properly normalized. The contaminating prior, $p_c(\theta)$ need not be in the same parametric family as $p_0(\theta)$, so as $p_c(\theta)$ ranges over all possible priors, equation (1) represents an expressive class of perturbations. Under mild assumptions (given in section §A), the local sensitivity measure at the prior $p(\theta|\epsilon)$ given by a particular ϵ is

$$S_\epsilon^{p_c} = \left. \frac{d\mathbb{E}_{p_\epsilon^x}[g(\theta)]}{d\epsilon} \right|_\epsilon = \text{Cov}_{p_\epsilon^x} \left(g(\theta), \frac{p_c(\theta) - p_0(\theta)}{p_0(\theta) + \epsilon(p_c(\theta) - p_0(\theta))} \right). \quad (2)$$

The definition in equation (2) depends on a choice of $p_c(\theta)$, which we denote with a superscript on $S_\epsilon^{p_c}$. At $\epsilon = 0$, we recover the local sensitivity around $p_0(\theta)$, which we denote $S_0^{p_c}$.

Rather than choose some finite set of $p_c(\theta)$ and calculate their corresponding $S_\epsilon^{p_c}$, one can work with a single function that summarizes the effect of any $p_c(\theta)$, called the ‘‘influence function’’ (Gustafson, 2000). Observing that equation (2) is a linear functional of $p_c(\theta)$ when $g(\theta)$, ϵ , and $p(\theta|\epsilon)$ are fixed, the influence function (when it exists) is defined as the linear operator $I_\epsilon(\theta)$ that characterizes the dependence of $S_\epsilon^{p_c}$ on $p_c(\theta)$:

$$S_\epsilon^{p_c} = \int I_\epsilon(\theta) p_c(\theta) d\theta \quad \text{where} \quad I_\epsilon(\theta) := \frac{p_\epsilon^x(\theta)}{p(\theta|\epsilon)} (g(\theta) - \mathbb{E}_{p_\epsilon^x}[g(\theta)]). \quad (3)$$

At $\epsilon = 0$, we recover the local sensitivity around $p_0(\theta)$, which we denote $I_0(\theta)$. When perturbing a low-dimensional marginal of the prior, $I_0(\theta)$ is an easy-to-visualize summary of the effect of sensitivity to an arbitrary $p_c(\theta)$ using quantities calculated only under $p_0(\theta)$ (see the example in section §4 and the extended discussion in Gustafson (2000)). Additionally, the worst case prior in a suitably defined metric ball around $p_0(\theta)$ is a functional of the influence function, as shown in Gustafson (2000).

2 Variational approximation and linear response

We now derive a version of equation (2) for Variational Bayes (VB) approximations to the posterior. Recall that a variational approximate posterior is a distribution selected to minimize the Kullback-Liebler (KL) divergence to p_ϵ^x across distributions q in some class \mathcal{Q} . Let q_ϵ^x denote the variational approximation to posterior p_ϵ^x . We assume that distributions in \mathcal{Q} are smoothly parameterized by a finite-dimensional parameter η whose optimum lies in the interior of some feasible set Ω_η .

We would like to calculate the local robustness measures of section §1 for the variational approximation q_ϵ^x , but a direct evaluation of the covariance in equation (2) can be misleading. For example, a common choice of the approximating family \mathcal{Q} is the class of distributions that factorize across θ . This is known as the ‘‘mean field approximation’’ (Wainwright and Jordan, 2008). By construction, a mean field approximation does not model covariances between independent components of θ , so a naive estimate of the covariance in equation (2) may erroneously suggest that the prior on one component of θ cannot affect the posterior on another.

However, for VB approximations, we can evaluate the derivative on the left hand side of equation (2) directly. Using linear response variational Bayes (LRVB) (Giordano et al., 2016, 2015), we have

$$\left. \frac{d}{d\epsilon} \mathbb{E}_{q_\epsilon^x}[g(\theta)] \right|_\epsilon = \int \frac{q_\epsilon^x(\theta)}{p(\theta|\epsilon)} q_\eta(\theta)^T H^{-1} g_\eta p_c(\theta) d\theta \quad (4)$$

where $g_\eta := \frac{\partial \mathbb{E}_{q_\epsilon^x}[g(\theta)]}{\partial \eta}$, $q_\eta(\theta) := \frac{\partial \log q_\epsilon^x(\theta; \eta)}{\partial \eta}$, and $H := \frac{\partial^2 KL(q_\epsilon^x(\theta; \eta) || p_\epsilon^x)}{\partial \eta \partial \eta^T}$.

It follows immediately from the definition in equation (3) that we can define the *variational influence function*

$$I_\epsilon^q(\theta) := \frac{q_\epsilon^x(\theta)}{p(\theta|\epsilon)} q_\eta(\theta)^T H^{-1} g_\eta \quad (5)$$

that captures the sensitivity of $\mathbb{E}_{q_\epsilon^x} [g(\theta)]$ just as $I_\epsilon(\theta)$ captures the sensitivity of $\mathbb{E}_{p_\epsilon^x} [g(\theta)]$.

The VB versions of epsilon sensitivity measures have some advantages and disadvantages relative to using Markov Chain Monte Carlo (MCMC) to evaluate the exact sensitivities in section §1. Using MCMC samples from p_0^x , one can form a Monte Carlo estimate of the covariance in equation (2), though the sample variance may be infinite when p_c has heavier tails than p_0 . The extent to which this is a real problem in practice will vary. Similarly, one must take care in numerically evaluating equation (4), since naively sampling from q_ϵ^x may also result in infinite variance due to the term $p(\theta|\epsilon)$ in the denominator. Since we have a closed form for q_ϵ^x , we can instead evaluate equation (4) as an integral over p_c using importance sampling, as described in section §D. Still, providing efficient estimates of equation (4) for high-dimensional, non-conjugate, heavy-tailed p_c remains a challenge. Finally, in contrast to equation (3), where we do not generally have a closed form expression for p_ϵ^x , every term in equation (5) is known. This means it is easier to evaluate the influence function for VB approximations than from MCMC draws, especially far from the posterior.

3 Non-local approximation

Equation (3) quantifies the effect of adding an infinitesimal amount of the contaminating prior. In practice, we may also want to evaluate intermediate values of ϵ , particularly $\epsilon = 1$, which represents completely replacing $p_0(\theta)$ with $p_c(\theta)$. Since $p_c(\theta)$ may be quite different from $p_0(\theta)$, this is a non-local robustness measure. For MCMC samples, one can use importance sampling, which is essentially equivalent to evaluating the covariance in equation (2) (with the same problem of infinite variance – see section §C). For VB, however, we either need to re-fit the model for each new prior (which may be time consuming) or somehow use the local information at q_0^x . In this paper, we investigate the latter. For the remainder of this section, since our results are general, we will discuss using local information in p_0^x . However, the reader should keep in mind that the ultimate goal is to apply the insights gained to the variational approximation q_0^x .

One might hope to linearly extrapolate from $\epsilon = 0$ to $\epsilon = 1$ using the slope $S_0^{p_c}$ at $\epsilon = 0$. That is, we might hope that $\mathbb{E}_{p_\epsilon^x} [g(\theta)] \Big|_{\epsilon=0}^{\epsilon=1} \approx (1 - 0) \frac{d\mathbb{E}_{p_\epsilon^x} [g(\theta)]}{d\epsilon} \Big|_{\epsilon=0}$. However, as we will now show, this is not realistic when one of the two priors is more consistent with the data than the other. Inspection of equation (3) shows that posterior expectations are highly sensitive to perturbations of priors which are inconsistent with the data: if $p(\theta|\epsilon)$ is small in an area of the θ space where p_ϵ^x is not small, then the influence function $I_\epsilon(\theta)$ will be quite large. The model will have high sensitivity to any contaminating prior, $p_c(\theta)$, that is more consistent with the model than $p(\theta|\epsilon)$ at ϵ . In particular, this is true at $\epsilon = 0$ where $p(\theta|\epsilon) = p_0(\theta)$ if $p_0(\theta)$ is inconsistent with the data. In fact, as we show in section §A,

$$\mathbb{E}_{p_\epsilon^x} [g(\theta)] \Big|_{\epsilon=0}^{\epsilon=1} = \int_0^1 \frac{d\mathbb{E}_{p_\epsilon^x} [g(\theta)]}{d\epsilon} \Big|_\epsilon d\epsilon = \frac{\int p(x|\theta) p_c(\theta) d\theta}{\int p(x|\theta) p_0(\theta) d\theta} S_0^{p_c}. \quad (6)$$

When the model evidence is very different for p_c and p_0 , e.g. when $\int p(x|\theta) p_0(\theta) d\theta \ll \int p(x|\theta) p_c(\theta) d\theta$ as in section §4, the extrapolated slope $S_0^{p_c}$ can be quite different from the effect of replacing completely replacing $p_0(\theta)$ with $p_c(\theta)$.

However, as ϵ grows away from zero and the new prior $p_c(\theta)$ is taken into account, the influence function will shrink. Observe that as a function of ϵ equation (2), one can show (see section §A) that

$$\frac{d\mathbb{E}_{p_\epsilon^x} [g(\theta)]}{d\epsilon} \Big|_\epsilon \leq \max \left\{ \frac{1}{\epsilon}, \frac{1}{1 - \epsilon} \right\} \mathbb{E}_{p_\epsilon^x} [|g(\theta) - \mathbb{E}_{p_\epsilon^x} [g(\theta)]|]. \quad (7)$$

For $\epsilon = 0$ or 1, this bound is vacuous, since the ratio $p_c(\theta) / p_0(\theta)$ can be arbitrarily large in areas assigned positive probability by p_ϵ^x . However, for intermediate values of ϵ , such as $\frac{1}{2}$, the bound is quite strong. In other words, contamination with $p_c(\theta)$ can have great influence on $\mathbb{E}_{p_\epsilon^x} [g(\theta)]$ when ϵ is close to the boundaries of $[0, 1]$, but once $p_c(\theta)$ is taken into account with intermediate ϵ , its influence is tightly bounded by equation (7). In this sense, the value of $\frac{d\mathbb{E}_{p_\epsilon^x} [g(\theta)]}{d\epsilon} \Big|_\epsilon$ is most atypical of its value across the interval $\epsilon \in [0, 1]$ at its endpoints. A real-life example of exactly this phenomenon is shown in section §4 in Fig. (1).

This suggests replacing the derivative at $\epsilon = 0$ with an average of the derivative over the interval $\epsilon \in [0, 1]$. To do this, note that the difficulty of the integral in equation (6) is the complicated dependence of p_ϵ^x on ϵ in $I_\epsilon(\theta)$. However, we can approximate the integral by keeping p_ϵ^x fixed at p_0^x so that $I_\epsilon(\theta)$ only depends on ϵ through $p(\theta|\epsilon)$. Under this approximation, the integral can be evaluated analytically (see section §B), giving the contaminating pseudo-density, $p_{mv}(\theta)$, which represents the approximate effect of integrating over ϵ from 0 to 1:

$$\mathbb{E}_{p_\epsilon^x} [g(\theta)] \Big|_{\epsilon=0}^{\epsilon=1} \approx \int I_0(\theta) p_{mv}(\theta) d\theta \quad p_{mv}(\theta) := \frac{p_c(\theta) p_0(\theta)}{p_c(\theta) - p_0(\theta)} \log \frac{p_c(\theta)}{p_0(\theta)}. \quad (8)$$

In the notation $p_{mv}(\theta)$, “mv” stands for “mean value”, by analogy with the mean value theorem for functions of real numbers. As shown in section §4, using $p_{mv}(\theta)$ with $I_0(\theta)$ rather than $p_c(\theta)$ can represent a significant improvement over equation (3) in practice when extrapolating to $\epsilon = 1$. We will focus on using p_{mv} with the variational approximations described in section §2.

4 Experiments

We demonstrate our methods using simulated data from a hierarchical model described in section §E. Here, we will demonstrate that our sensitivity measures accurately predict the changes in VB solutions. We discuss the close match between the VB and MCMC results in section §E. The results below are for the sensitivity of the expectation $\mathbb{E}_{p_0^x} [\mu_{11}]$ to the prior $p(\mu)$, though similar results are easily calculated for any other low-dimensional prior marginal or posterior expectation.

We generate data using true parameters that are far from $p_0(\theta)$ so that our model is not robust to perturbations, as can be seen by the large values of the influence function, which is pictured in the top left panel of Fig. (1). The posterior mean is shown with a black dot, indicating that, though large, the influence function is very highly concentrated around the posterior mean. The top right panel of Fig. (1) indicates how $\mathbb{E}_{q_\epsilon^x} [\mu_{11}]$ depends on ϵ near $\epsilon = 0$. The slope is very steep at $\epsilon = 0$, reflecting the fact that $p_c(\theta)$ takes values much larger than $p_0(\theta)$ near the posterior where the influence function is very high. However, it very quickly levels off for ϵ only slightly above zero.

The top right panel of Fig. (1) indicates that extrapolating from the slope $S_0^{p_c}$ at $\epsilon = 0$ will radically over-estimate the effect of replacing $p_0(\theta)$ with $p_c(\theta)$. This is confirmed in the bottom left panel, which has the actual change in $\mathbb{E}_{q_\epsilon^x} [\theta]$ on the x-axis and $\int I_\epsilon^q(\theta) p_c(\theta) d\theta$ on the y-axis. Clearly, the extrapolation is unrealistic. However, the right panel of Fig. (1) demonstrates that $\int I_\epsilon^q(\theta) p_{mv}(\theta) d\theta$ accurately matches the effect of replacing $p_0(\theta)$ with $p_c(\theta)$. Note the different ranges in the y-axis (which prohibit plotting the two graphs on the same scale). The error bars represent importance sampling error.

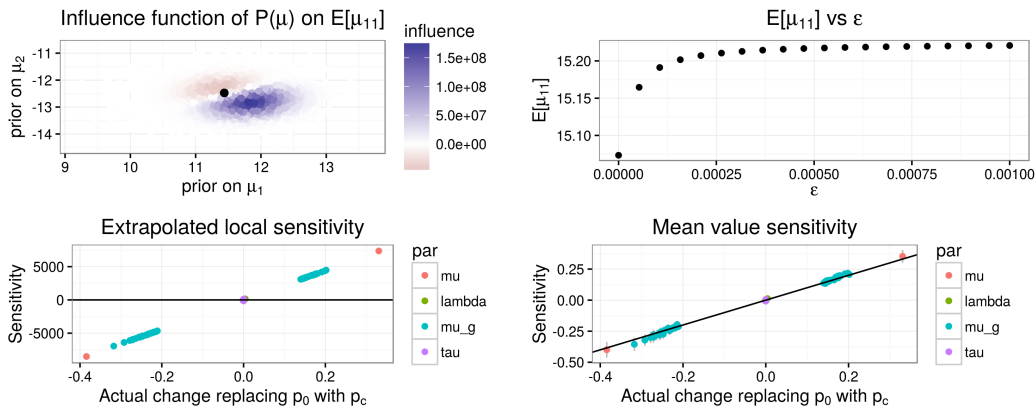


Figure 1: Simulation results

Acknowledgements The paper was greatly improved through the insightful comments of our anonymous reviewers. Ryan Giordano was supported by the Director, Office of Science, Office of Advanced Scientific Computing Research, Applied Mathematics program of the U.S. Department of Energy under Contract No. DE-AC02-05CH1123.

References

- John Barnard, Robert McCulloch, and Xiao-Li Meng. Modeling covariance matrices in terms of standard deviations and correlations, with application to shrinkage. *Statistica Sinica*, 10(4): 1281–1312, 2000.
- James O. Berger, David Ríos Insua, and Fabrizio Ruggeri. Robust Bayesian analysis. In David Ríos Insua and Fabrizio Ruggeri, editors, *Robust Bayesian Analysis*, volume 152. Springer Science & Business Media, 2000.
- Rick Durrett. *Probability: theory and examples*. Cambridge University Press, 2010.
- Ryan Giordano, Tamara Broderick, and Michael Jordan. Linear response methods for accurate covariance estimates from mean field variational Bayes. *arXiv preprint arXiv:1506.04088*, 2015.
- Ryan Giordano, Tamara Broderick, Rachael Meager, Jonathan Huggins, and Michael Jordan. Fast robustness quantification with variational Bayes. *arXiv preprint arXiv:1606.07153*, 2016.
- Paul Gustafson. Local sensitivity of posterior expectations. *The Annals of Statistics*, 24(1):174–195, 1996.
- Paul Gustafson. Local robustness in Bayesian analysis. In David Ríos Insua and Fabrizio Ruggeri, editors, *Robust Bayesian Analysis*, volume 152. Springer Science & Business Media, 2000.
- Daniel Lewandowski, Dorota Kurowicka, and Harry Joe. Generating random correlation matrices based on vines and extended onion method. *Journal of multivariate analysis*, 100(9):1989–2001, 2009.
- Rachael Meager. Understanding the impact of microcredit expansions: A Bayesian hierarchical analysis of 7 randomised experiments. *arXiv preprint arXiv:1506.06669*, 2015.
- Donald B Rubin. Estimation in parallel randomized experiments. *Journal of Educational and Behavioral Statistics*, 6(4):377–401, 1981.
- Stan Team. *Stan Modeling Language Users Guide and Reference Manual, Version 2.8.0*, 2015. URL <http://mc-stan.org/>.
- M. J. Wainwright and M. I. Jordan. Graphical models, exponential families, and variational inference. *Foundations and Trends® in Machine Learning*, 1(1-2):1–305, 2008.

Appendices

A Epsilon sensitivity

Throughout the paper, we make the following assumptions:

- Assumption 1. For all $\epsilon \in [0, 1]$, $p(\theta|\epsilon)$ is strictly positive where θ has positive measure.
- Assumption 2. Both $p(\theta|x)$ and $g(\theta)p(\theta|x)$ are bounded as a function of θ .

Under these assumptions, equation (2) follows directly from Gustafson (1996) result 8. For completeness, we reproduce a slightly simpler proof under the slightly less clearly-articulated assumption that we can exchange integration and differentiation as described, for example, in Durrett (2010, Appendix A.5). Denote by $p(\theta|\epsilon)$ any distribution parameterized by the scalar ϵ (not necessarily a prior). Then by direct differentiation,

$$\begin{aligned} \frac{d\mathbb{E}_{p(\theta|\epsilon)}[g(\theta)]}{d\epsilon} &= \frac{d}{d\epsilon} \int p(\theta|\epsilon) g(\theta) d\theta \\ &= \int \frac{dp(\theta|\epsilon)}{d\epsilon} g(\theta) d\theta \\ &= \int \frac{d \log p(\theta|\epsilon)}{d\epsilon} p(\theta|\epsilon) g(\theta) d\theta. \end{aligned} \quad (9)$$

By applying equation (9) to $g(\theta) = 1$, we see that $\mathbb{E}_{p(\theta|\epsilon)} \left[\frac{d \log p(\theta|\epsilon)}{d\epsilon} \right] = \int \frac{d \log p(\theta|\epsilon)}{d\epsilon} p(\theta|\epsilon) d\theta = 0$, so we can subtract $0 = \mathbb{E}_{p(\theta|\epsilon)}[g(\theta)] \mathbb{E}_{p(\theta|\epsilon)} \left[\frac{d \log p(\theta|\epsilon)}{d\epsilon} \right]$ to get

$$\frac{d\mathbb{E}_{p(\theta|\epsilon)}[g(\theta)]}{d\epsilon} = \text{Cov}_{p(\theta|\epsilon)} \left(g(\theta), \frac{d \log p(\theta|\epsilon)}{d\epsilon} \right). \quad (10)$$

To derive equation (2), we simply observe that

$$\begin{aligned} \frac{d \log p(\theta|\epsilon)}{d\epsilon} &= \frac{d}{d\epsilon} \log((1-\epsilon)p_0(\theta) + \epsilon p_c(\theta)) \\ &= \frac{p_c(\theta) - p_0(\theta)}{p_0(\theta) + \epsilon(p_c(\theta) - p_0(\theta))}. \end{aligned}$$

Note that the assumptions also suffice to assure that the covariance is bounded.

Next, we observe the simple relationship between epsilon sensitivity at $\epsilon = 0$ and the effect of replacing one prior with another. First, defining the normalizing constants

$$\begin{aligned} C_0 &:= \int p(x|\theta) p_0(\theta) d\theta \\ C_1 &:= \int p(x|\theta) p_c(\theta) d\theta \\ \mathbb{E}_{p_0^x} \left[\frac{p_c(\theta)}{p_0(\theta)} \right] &= \int \frac{p(x|\theta) p_0(\theta) p_c(\theta)}{C_0 p_0(\theta)} d\theta = \frac{C_1}{C_0}, \end{aligned}$$

by straightforward manipulation we have

$$\begin{aligned}
\mathbb{E}_{p_\epsilon^x(\theta)} [g(\theta)] \Big|_{\epsilon=1} - \mathbb{E}_{p_\epsilon^x(\theta)} [g(\theta)] \Big|_{\epsilon=0} &= \int \frac{p(x|\theta) p_c(\theta)}{C_1} g(\theta) d\theta - \mathbb{E}_{p_0^x(\theta)} [g(\theta)] \\
&= \int \frac{p(x|\theta) p_c(\theta) p_0(\theta)}{C_1 C_0} \frac{C_0}{p_0(\theta)} g(\theta) d\theta - \frac{C_0}{C_1} \mathbb{E}_{p_0^x} \left[\frac{p_c(\theta)}{p_0(\theta)} \right] \mathbb{E}_{p_0^x(\theta)} [g(\theta)] \\
&= \frac{C_0}{C_1} \left(\mathbb{E}_{p_0^x(\theta)} \left[\frac{p_c(\theta)}{p_0(\theta)} g(\theta) \right] - \mathbb{E}_{p_0^x} \left[\frac{p_c(\theta)}{p_0(\theta)} \right] \mathbb{E}_{p_0^x(\theta)} [g(\theta)] \right) \\
&= \frac{C_0}{C_1} \text{Cov}_{p_0^x(\theta)} \left(g(\theta), \frac{p_c(\theta)}{p_0(\theta)} \right) \\
&= \frac{C_0}{C_1} \frac{d\mathbb{E}_{p_\epsilon^x} [g(\theta)]}{d\epsilon} \Big|_{\epsilon=0}.
\end{aligned}$$

Finally, we derive equation (7) using equation (2):

$$\begin{aligned}
\left| \frac{d\mathbb{E}_{p_\epsilon^x} [g(\theta)]}{d\epsilon} \right| \Big|_\epsilon &= \left| \text{Cov}_{p_\epsilon^x} \left(g(\theta), \frac{p_c(\theta) - p_0(\theta)}{p_0(\theta) + \epsilon(p_c(\theta) - p_0(\theta))} \right) \right| \\
&\leq \mathbb{E}_{p_\epsilon^x} [|g(\theta) - \mathbb{E}_{p_\epsilon^x} [g(\theta)]|] B_\epsilon
\end{aligned}$$

where

$$\begin{aligned}
B_\epsilon &= \sup_\theta \left| \frac{p_c(\theta) - p_0(\theta)}{p_0(\theta) + \epsilon(p_c(\theta) - p_0(\theta))} \right| \\
&= \sup_\theta \begin{cases} \frac{p_c(\theta) - p_0(\theta)}{p_0(\theta) + \epsilon(p_c(\theta) - p_0(\theta))} & \text{when } p_c(\theta) > p_0(\theta) \\ \frac{p_0(\theta) - p_c(\theta)}{p_0(\theta) + \epsilon(p_c(\theta) - p_0(\theta))} & \text{when } p_c(\theta) \leq p_0(\theta) \end{cases} \\
&= \sup_\theta \begin{cases} \frac{p_c(\theta) - p_0(\theta)}{p_0(\theta) + \epsilon(p_c(\theta) - p_0(\theta))} & \text{when } p_c(\theta) > p_0(\theta) \\ \frac{p_0(\theta) - p_c(\theta)}{p_c(\theta) + (1-\epsilon)(p_0(\theta) - p_c(\theta))} & \text{when } p_c(\theta) \leq p_0(\theta) \end{cases} \\
&\leq \sup_\theta \begin{cases} \frac{p_c(\theta) - p_0(\theta)}{\epsilon(p_c(\theta) - p_0(\theta))} & \text{when } p_c(\theta) > p_0(\theta) \\ \frac{p_0(\theta) - p_c(\theta)}{(1-\epsilon)(p_0(\theta) - p_c(\theta))} & \text{when } p_c(\theta) \leq p_0(\theta) \end{cases} \\
&\leq \max \left\{ \frac{1}{\epsilon}, \frac{1}{1-\epsilon} \right\}.
\end{aligned}$$

B Mean Value Contaminating Prior

Under the assumption that $p_\epsilon^x \approx p_0^x$,

$$\begin{aligned}
\mathbb{E}_{p_\epsilon^x} [g(\theta)] \Big|_{\epsilon=1} - \mathbb{E}_{p_\epsilon^x} [g(\theta)] \Big|_{\epsilon=0} &= \int_0^1 \int \frac{p_\epsilon^x(\theta) p_c(\theta)}{p(\theta|\epsilon)} \left(g(\theta) - \mathbb{E}_{p_\epsilon^x(\theta)|\epsilon=0} [g(\theta)] \right) d\theta d\epsilon \\
&\approx \int_0^1 \int \frac{p_0^x(\theta) p_c(\theta)}{p(\theta|\epsilon)} \left(g(\theta) - \mathbb{E}_{p_0^x(\theta)|\epsilon=0} [g(\theta)] \right) d\theta d\epsilon \\
&= \int p_0^x(\theta) \left(g(\theta) - \mathbb{E}_{p_0^x(\theta)} [g(\theta)] \right) \int_0^1 \frac{p_c(\theta)}{p(\theta|\epsilon)} d\epsilon d\theta \\
&= \int p_0^x(\theta) \left(g(\theta) - \mathbb{E}_{p_0^x(\theta)} [g(\theta)] \right) \frac{p_c(\theta)}{p_c(\theta) - p_0(\theta)} (\log p_c(\theta) - \log p_0(\theta)) d\theta \\
&= \int I_0(\theta) \frac{p_c(\theta) p_0(\theta)}{p_c(\theta) - p_0(\theta)} (\log p_c(\theta) - \log p_0(\theta)) d\theta.
\end{aligned}$$

Where we have used

$$\begin{aligned}
\int_0^1 \frac{p_c(\theta)}{p(\theta|\epsilon)} d\epsilon &= \int_0^1 \frac{p_c(\theta)}{(1-\epsilon)p_0(\theta) + \epsilon p_c(\theta)} d\epsilon \\
&= \frac{p_c(\theta)}{p_c(\theta) - p_0(\theta)} \int_0^1 \frac{d}{d\epsilon} \log((1-\epsilon)p_0(\theta) + \epsilon p_c(\theta)) d\epsilon \\
&= \frac{p_c(\theta)}{p_c(\theta) - p_0(\theta)} (\log p_c(\theta) - \log p_0(\theta)).
\end{aligned}$$

Consequently, applying equation (3) with the pseudo-density

$$p_{mv} := \frac{p_c(\theta) p_0(\theta)}{p_c(\theta) - p_0(\theta)} (\log p_c(\theta) - \log p_0(\theta))$$

represents an approximation to the quantity $\mathbb{E}_{p_\epsilon^x} [g(\theta)]|_{\epsilon=1} - \mathbb{E}_{p_\epsilon^x} [g(\theta)]|_{\epsilon=0}$, which is the effect of completely replacing $p_0(\theta)$ with $p_c(\theta)$.

C Comparison with MCMC importance sampling

In this section, we show that using importance sampling with MCMC samples to calculate the local sensitivity equation (2) is precisely equivalent to using the same MCMC samples to estimate the covariance in equation (10) directly. Suppose, without loss of generality, we have samples θ_i drawn from $p_0^x(\theta)$

$$\begin{aligned}
\theta_i &\sim p_0^x(\theta) \\
\mathbb{E}_{p_0^x} [g(\theta)] &\approx \sum_i g(\theta_i).
\end{aligned}$$

If we could calculate the normalizing constants, the importance sampling estimate for $\mathbb{E}_{p_\epsilon^x} [g(\theta)]$ would be

$$\begin{aligned}
\mathbb{E}_{p_\epsilon^x} [g(\theta)] &\approx \frac{1}{N} \sum_i w_i g(\theta_i) \\
w_i &:= \frac{p_\epsilon^x(\theta)}{p_0^x(\theta)} = \frac{p(x|\theta) p(\theta|\epsilon)}{C_\epsilon} \frac{C_0}{p(x|\theta) p_0(\theta)} \\
&= \exp(\log p(\theta|\epsilon) - \log p_0(\theta) + \log C_0 - \log C_\epsilon) \\
C_0 &:= \int p(x|\theta) p_0(\theta) d\theta \\
C_\epsilon &:= \int p(x|\theta) p(\theta|\epsilon) d\theta.
\end{aligned}$$

Differentiating the weights,

$$\begin{aligned}
\frac{dw_i}{d\epsilon} &= w_i \left(\frac{d \log p(\theta|\epsilon)}{d\epsilon} - \frac{d \log C_\epsilon}{d\epsilon} \right) \\
&= w_i \left(\frac{d \log p(\theta|\epsilon)}{d\epsilon} - \frac{1}{C_\epsilon} \int p(x|\theta) p(\theta|\epsilon) \frac{d \log p(\theta|\epsilon)}{d\epsilon} d\theta \right) \\
&= w_i \left(\frac{d \log p(\theta|\epsilon)}{d\epsilon} - \mathbb{E}_{p_\epsilon^x} \left[\frac{d \log p(\theta|\epsilon)}{d\epsilon} \right] \right).
\end{aligned}$$

It follows that

$$\frac{d}{d\epsilon} \frac{1}{N} \sum_i w_i g(\theta_i) \Big|_\epsilon = \frac{1}{N} \sum_i w_i \left(\frac{d \log p(\theta_i|\epsilon)}{d\epsilon} - \mathbb{E}_{p_\epsilon^x} \left[\frac{d \log p(\theta|\epsilon)}{d\epsilon} \right] \right) g(\theta_i)$$

which is precisely the MCMC estimate of the covariance given by equation (2). In particular, when $\epsilon = 0$, we have

$$\left. \frac{d}{d\epsilon} \frac{1}{N} \sum_i w_i g(\theta_i) \right|_{\epsilon=0} = \frac{1}{N} \sum_i \left(\frac{p_c(\theta_i)}{p_0(\theta_i)} - \mathbb{E}_{p_\epsilon} \left[\frac{p_c(\theta)}{p_0(\theta)} \right] \right) g(\theta_i)$$

and the importance sampling estimate for replacing p_0 with p_c is

$$\begin{aligned} \left. \sum_i w_i g(\theta_i) \right|_{\epsilon=1} &= \frac{1}{N} \sum_i \left(\frac{C_0 p_c(\theta_i)}{C_1 p_0(\theta_i)} - \mathbb{E}_{p_\epsilon} \left[\frac{p_c(\theta_i)}{p_0(\theta_i)} \right] \right) g(\theta_i) \\ &= \frac{C_0}{C_1} \left. \frac{d}{d\epsilon} \frac{1}{N} \sum_i w_i g(\theta_i) \right|_{\epsilon=0} \end{aligned}$$

which confirms that the importance sampling estimate is exactly the Monte Carlo analogue of the relation equation (6).

In general, we do not know C_0 and C_ϵ and must use instead

$$\begin{aligned} \omega_i &:= \frac{p(\theta|\epsilon)}{p_0(\theta)} \\ \tilde{\omega}_i &:= \frac{\omega_i}{\sum_j \omega_j}. \end{aligned}$$

Then

$$\begin{aligned} \frac{d\tilde{\omega}_i}{d\epsilon} &= \frac{\omega_i}{\sum_j \omega_j} \frac{d \log p(\theta_i|\epsilon)}{d\epsilon} - \frac{\omega_i}{(\sum_k \omega_k)^2} \sum_j \omega_j \frac{d \log p(\theta_j|\epsilon)}{d\epsilon} \\ &= \frac{\omega_i}{\sum_j \omega_j} \left(\frac{d \log p(\theta_i|\epsilon)}{d\epsilon} - \sum_j \frac{\omega_j}{\sum_k \omega_k} \frac{d \log p(\theta_j|\epsilon)}{d\epsilon} \right) \end{aligned}$$

which simply replaces the (possibly intractable) expectation $\mathbb{E}_{p_\epsilon} \left[\frac{d \log p(\theta|\epsilon)}{d\epsilon} \right]$ with its MCMC estimate.

D Variational Bayes importance sampling

To evaluate equation (4) requires approximate integration for which we use importance sampling:

$$\begin{aligned} \theta_s &\sim u(\theta), \text{ for } s = 1 : S \\ w_s &:= \frac{p_c(\theta_s)}{u(\theta_s)} \\ \int \frac{q_0^x(\theta)}{p_0(\theta)} q_\eta(\theta)^T H^{-1} g_\eta p_c(\theta) d\theta &\approx \frac{1}{S} \sum_{s=1}^S \frac{q_0^x(\theta_s)}{p_0(\theta_s)} q_\eta(\theta_s)^T H^{-1} g_\eta w_s \\ &= \frac{1}{S} \sum_{s=1}^S I_0^q(\theta_s) w_s. \end{aligned}$$

Note that the influence function can be evaluated once for a large number of draws from $u(\theta)$, and then the weights and prior density can be quickly calculated for any perturbation $p_c(\theta)$, allowing for fast computation of sensitivity to any $p_c(\theta)$ with little additional overhead.

Since the influence function is mostly concentrated around $q_0^x(\theta)$, we set $u(\theta)$ to be $q_0^x(\theta)$ but with quadrupled variance (so that standard deviations are doubled). Note that this choice of u is a poor approximation of p_c , which is nominally the target distribution for importance sampling. However,

since $I_0^q(\theta)$ is very small far from $q_0^x(\theta)$, it is an adequate approximation of the integral equation (4). Formally, suppose that $I_0^q(\theta)$ is concentrated on a set A in the sense that $\sup_{\theta \in A^c} |I_0^q(\theta)| \leq \epsilon$ for some small δ . Then the absolute error in evaluating equation (4) on the set A only is also bounded by δ :

$$\begin{aligned} \left| \int I_0^q(\theta) p_c(\theta) d\theta - \int_A I_0^q(\theta) p_c(\theta) d\theta \right| &= \left| \int_{A^c} I_0^q(\theta) p_c(\theta) d\theta \right| \\ &\leq \int_{A^c} |I_0^q(\theta)| p_c(\theta) d\theta \\ &\leq \sup_{\theta \in A^c} |I_0^q(\theta)| \int_{A^c} p_c(\theta) d\theta \\ &\leq \delta. \end{aligned}$$

As long as $q_\epsilon^x(\theta)$ is chosen to have lighter tails than $p_0(\theta)$ (which is determined by \mathcal{Q}), $I_0^q(\theta)$ will decay quickly away from $q_0^x(\theta)$, and we can choose A centered on $q_0^x(\theta)$. Consequently, we can think of u as approximating $p_c(\theta) 1_A$ rather than $p_c(\theta)$.

E Microcredit model

We simulate data using a variant of the analysis performed in (Meager, 2015), though with somewhat different prior choices. In Meager (2015), randomized controlled trials were run in seven different sites to try to measure the effect of access to microcredit on various measures of business success. Each trial was found to lack power individually for various reasons, so there could be some benefit to pooling the results in a simple hierarchical model. For the purposes of demonstrating robust Bayes techniques with VB, we will focus on the simpler of the two models in (Meager, 2015) and ignore covariate information.

We will index sites with $k = 1, \dots, K$ (here, $K = 30$) and business within a site by $i = 1, \dots, N_k$. The total number of observations was $\sum_k N_k = 3000$. In site k and business i we observe whether the business was randomly selected for increased access to microcredit, denoted T_{ik} , and the profit after intervention, y_{ik} . We follow (Rubin, 1981) and assume that each site has an idiosyncratic average profit, μ_{k1} and average improvement in profit, μ_{k2} , due to the intervention. Given μ_k , τ_k , and T_{ik} , the _observed profit is assumed to be generated according to

$$\begin{aligned} y_{ik} | \mu_k, \tau_k, x_{ik}, \sigma_k &\sim N(\mu_k^T x_{ik}, \sigma_k^2) \\ x_{ik} &:= \begin{pmatrix} 1 \\ T_{ik} \end{pmatrix}. \end{aligned}$$

The site effects, μ_k , are assumed to come from an overall pool of effects and may be correlated:

$$\begin{aligned} \mu_k | \mu &\sim N(\mu, C) \\ C &:= \begin{pmatrix} \sigma_\mu^2 & \sigma_{\mu\tau} \\ \sigma_{\mu\tau} & \sigma_\tau^2 \end{pmatrix}. \end{aligned}$$

The effects μ and the covariance matrix V are unknown parameters that require priors. For μ we simply use a bivariate normal prior. However, choosing an appropriate prior for a covariance matrix can be conceptually difficult (Barnard et al., 2000). Following the recommended practice of the software package STAN (Stan Team, 2015), we derive a variational model to accommodate the non-conjugate LKJ prior (Lewandowski et al., 2009), allowing the user to model the covariance and marginal variances separately. Specifically, we use

$$\begin{aligned}
C &=: SRS \\
S &= \text{Diagonal matrix} \\
R &= \text{Covariance matrix} \\
S_{kk} &= \sqrt{\text{diag}(C)_k}
\end{aligned}$$

We can then put independent priors on the scale of the variances, S_{kk} , and on the covariance matrix, R . We model the inverse of C with a Wishart variational distribution, and use the following priors:

$$\begin{aligned}
q(C^{-1}) &= \text{Wishart}(V_\Lambda, n) \\
p_0(S) &= \prod_{k=1}^2 p(S_{kk}) \\
S_{kk}^2 &\sim \text{InverseGamma}(\alpha_{scale}, \beta_{scale}) \\
\log p_0(R) &= (\eta - 1) \log |R| + C
\end{aligned}$$

The necessary expectations have closed forms with the Wishart variational approximation, as derived in Giordano et al. (2016). In addition, we put a normal prior on $(\mu, \tau)^T$ and an inverse gamma prior on σ_k^2 :

$$\begin{aligned}
p_0(\mu) &= N(0, \Lambda^{-1}) \\
p_0(\sigma_k^2) &= \text{InverseGamma}(\alpha_\tau, \beta_\tau)
\end{aligned} \tag{11}$$

The prior parameters used were:

$$\begin{aligned}
\Lambda &= \begin{pmatrix} 0.111 & 0 \\ 0 & 0.111 \end{pmatrix} \\
\eta &= 15.010 \\
\sigma_k^{-2} &\sim \text{InverseGamma}(2.010, 2.010) \\
\alpha_{scale} &= 20.010 \\
\beta_{scale} &= 20.010 \\
\alpha_\tau &= 2.010 \\
\beta_\tau &= 2.010
\end{aligned}$$

As seen in Fig. (2), the means in VB and MCMC match closely.




Figure 2: Comparison with MCMC

In our simulation, μ and μ_k are in \mathbb{R}^2 , so the domain of the prior $p_0(\mu)$ is two-dimensional and $I_0^q(\theta)$ can be easily visualized. We consider the problem of estimating the effect of replacing the prior on μ with a product of independent centered Student t priors. In the notation of section §1, we take

$$p_0(\mu) = \mathcal{N}(\mu_1; 0, \Lambda^{-1}) \cdot \mathcal{N}(\mu_2; 0, \Lambda^{-1}) \quad p_c(\theta) = \text{Student}(\mu_1; \nu) \cdot \text{Student}(\mu_2; \nu).$$

We leave all other priors the same, i.e. $p_0(\tau_k) = p_c(\tau_k)$ and $p_0(C) = p_c(C)$. In our case, we used $\nu = 1$ and $\Lambda = 0.111$. We will present sensitivity of $\mathbb{E}_{q_\epsilon^\pi}[\mu_{11}]$, the first component of the first

top-level effect. In the notation of section §1, we are taking $g(\theta) = \mu_{11}$. Most of the computation is in generating draws and values for the importance sampling of the influence function, which can be done once and then reused for any choice of $p_c(\theta)$ and $g(\theta)$.