



Supplement of

Chemical sparsity in Bayesian receptor models for aerosol source apportionment

Marta Via et al.

Correspondence to: Marta Via (marta.viagonzalez@ung.si) and Kaspar R. Daellenbach (kaspar.daellenbach@psi.ch)

The copyright of individual parts of the supplement might differ from the article licence.

1. Supplementary figures

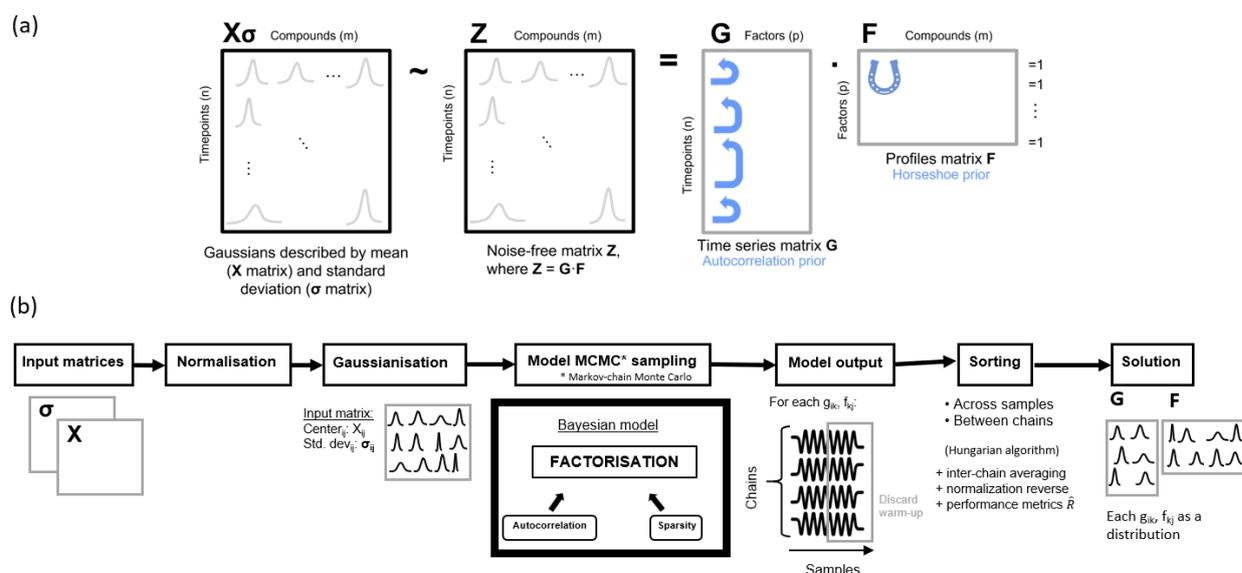
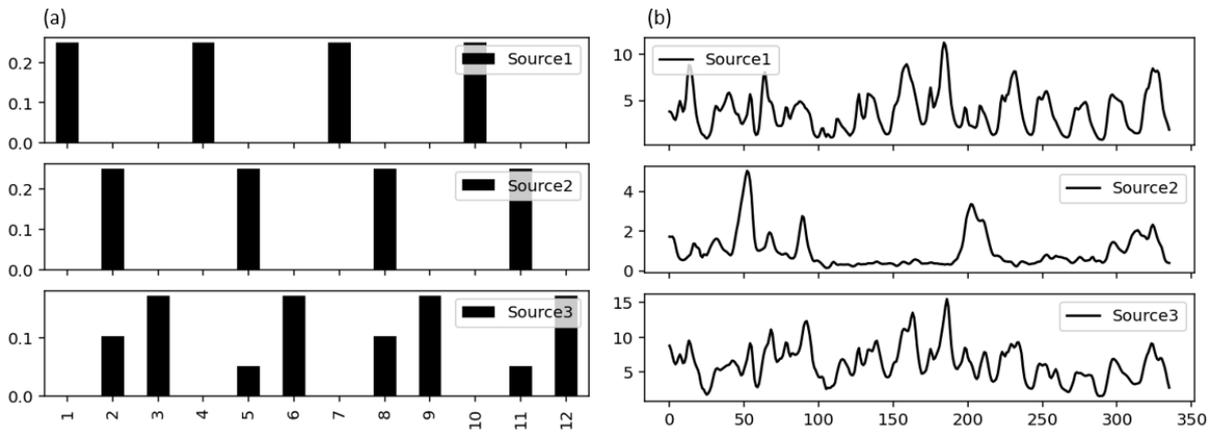


Figure S1. (a) Bayesian matrix factorisation for source apportionment sketch with autocorrelation and horseshoe priors. (b) Workflow diagram showing the aerosol source apportionment stages with the BAMF+HS model. This figure portrays the previous and posterior processes to the MCMC sampling, and both starting and end points. Concepts as chains and samples, mentioned in the text, are schematised in the workflow as well.

Table S1. Hamiltonian MonteCarlo sampling parameters for all the conducted experiments

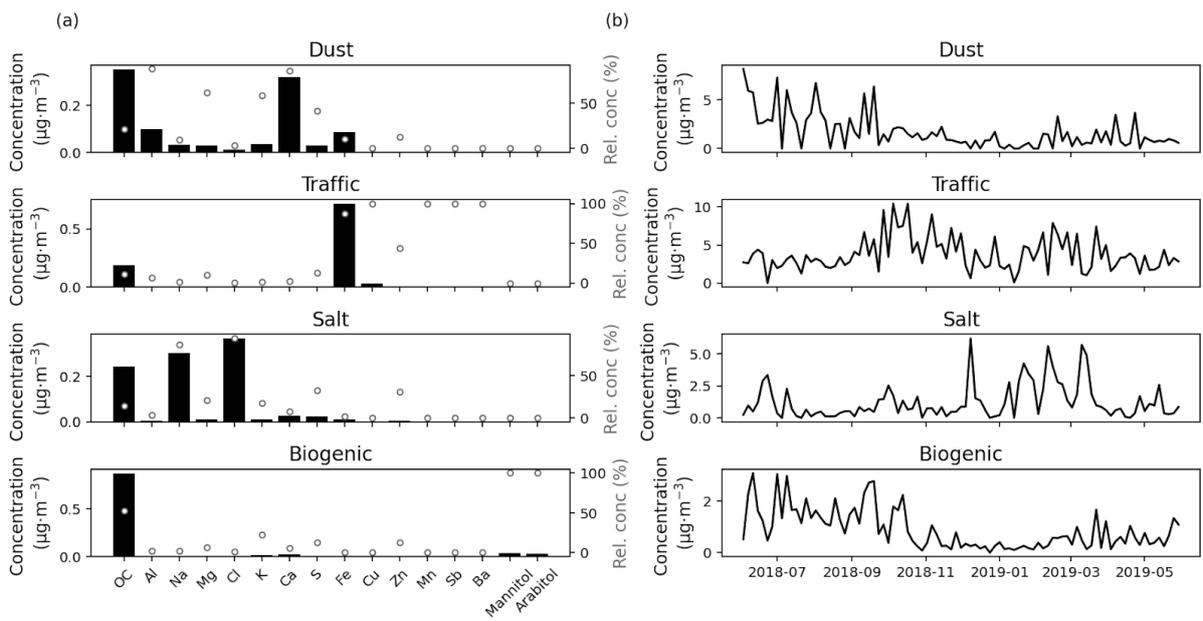
Experiment	Total number of samples	Number of warm-up samples	Number of chains	Time *
Toy dataset	4000	2000	4	~1h
European synthetic datasets	12000	6000	4	~7h
Filters synthetic dataset	6000	3000	4	~3h
Filters real-world dataset	6000	3000	4	~1.5h

* The time here shows the BAMF+HS running time plus the post-processing time (for sorting, averaging, etc.), the latter having a nearly neglectable contribution. Model runtimes were measured as wall-clock time on a high-performance computing cluster. Jobs were executed on compute nodes equipped with dual-socket AMD EPYC 7502 processors (64 physical cores, 2.5 GHz) and approximately 256 GB of RAM, running Linux. All BAMF+HS runs used four parallel chains using approximately 4 physical CPU cores (8 logical threads).



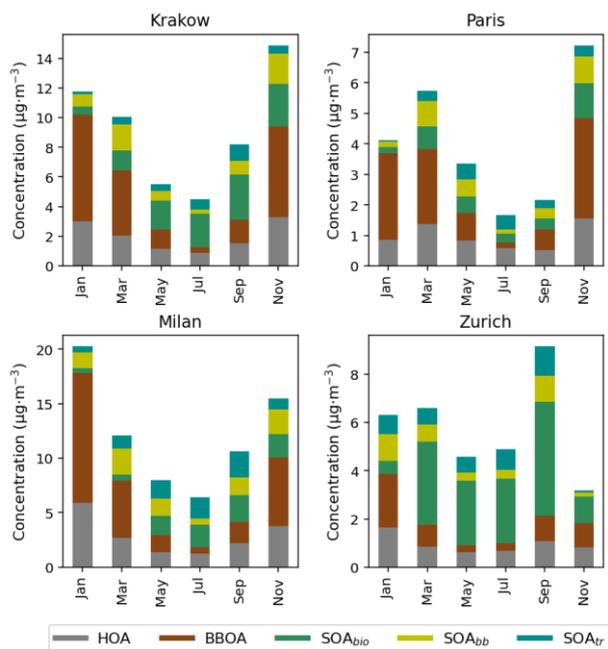
15

16 **Figure S2. Toy dataset (a) Profiles. (b) Time series.**



17

18 **Figure S3. Synthetic offline dataset (a) Profiles. (b) Time series.**



19

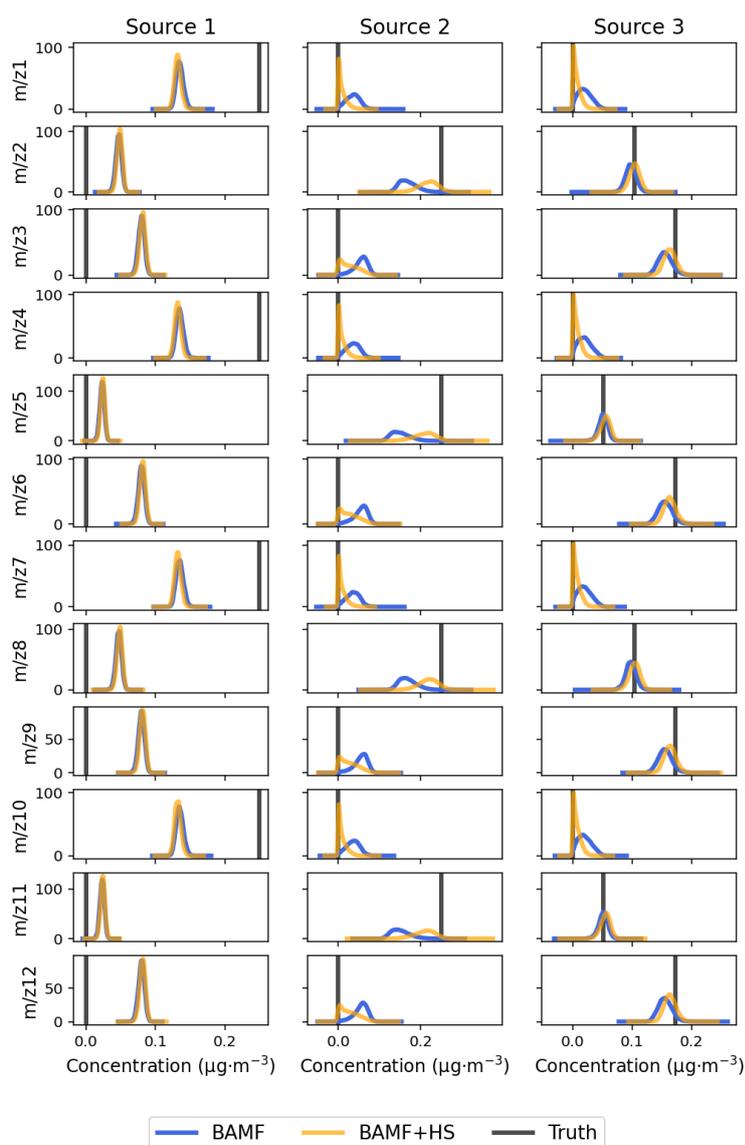
20 **Figure S4. Generated synthetic datasets source mean concentrations for 4 European cities.**

21

22 **Table S2. Profiles employed for the 4-cities synthetic datasets.**

City	Source	Citation
Krakow	HOA	Mohr et al. (2012)
	BBOA	Tobler et al. (2021)
Milan	HOA	Via et al. (2021)
	BBOA	Daellenbach et al. (2021)
Paris	HOA	Crippa et al. (2011)
	BBOA	Zhang et al. (2019)
Zurich	HOA	Elser et al. (2016)
	BBOA	Ulbrich et al. (2002), Ulbrich et al. (2022)
	SOA _{bio}	Daellenbach et al. (2017)
	SOA _{bb}	Ulbrich et al. (2002)
	SOA _{tr}	Sage et al. (2007)

23

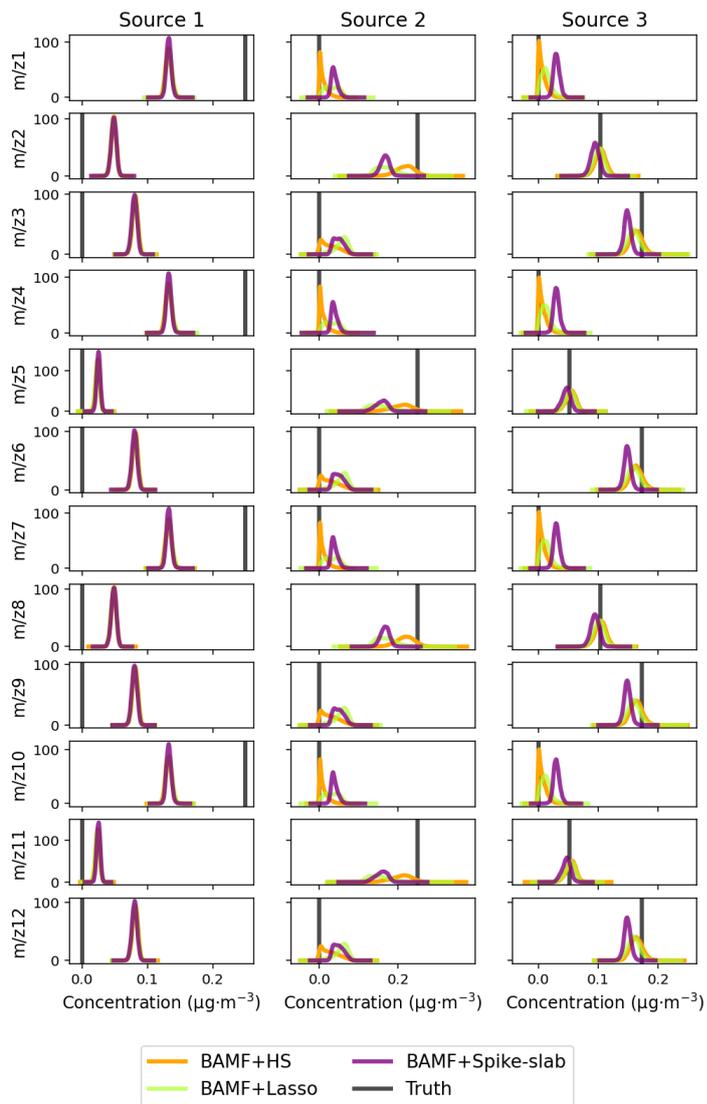


25

26 **Figure S5. Toy dataset distribution of F components for BAMF, BAMF+HS.**27 **Table S3. Statistics of the toy dataset reconstruction and comparison to truth for the different shrinkage alternatives.**

Model	Sources	X			G			F		
		R2	Median(Z-X /sigma)	Max(Z-X /sigma)	G/G ₀	R ²	ρ	R ²	Spars. ratio	Gini ratio
BAMF+HS	Source1	0.9689	0.2107	1.2863	1.9633	0.9850	0.8193	0.7551	0.0	0.40
	Source2				1.2786	0.9499	0.8193	0.9940	0.5	0.86
	Source3				0.5129	0.4781	0.9608	0.9994	1.0	0.93
BAMF-Lasso	Source1	0.9816	0.1239	0.6964	1.7579	0.9771	0.8193	0.7610	0.0	0.40
	Source2				2.0186	0.9491	0.8193	0.9515	0.0	0.50
	Source3				0.4910	0.4669	0.9608	0.9996	0.0	0.86
BAMF-Spike-Slab	Source1	0.9814	0.1262	0.7094	1.5725	0.9707	0.8193	0.7746	0.0	0.40
	Source2				1.7830	0.9490	0.8193	0.9936	0.0	0.52
	Source3				0.6215	0.4954	0.9608	0.9829	0.0	0.69

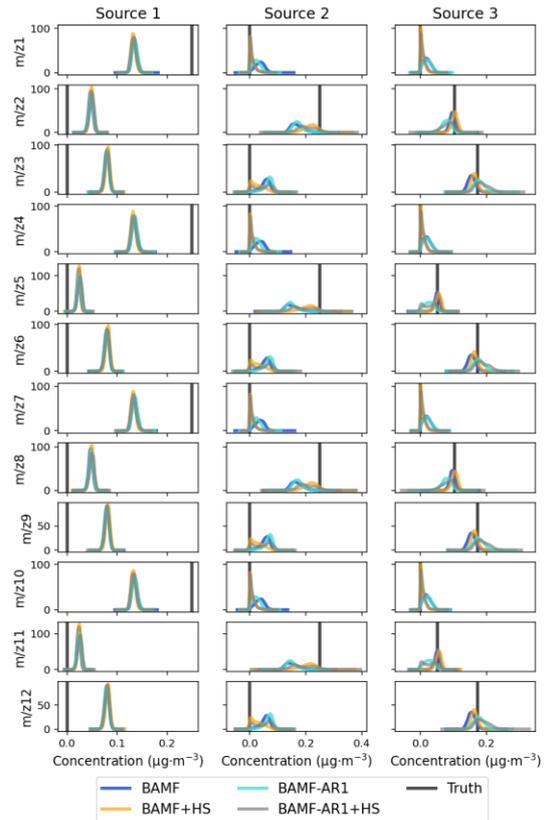
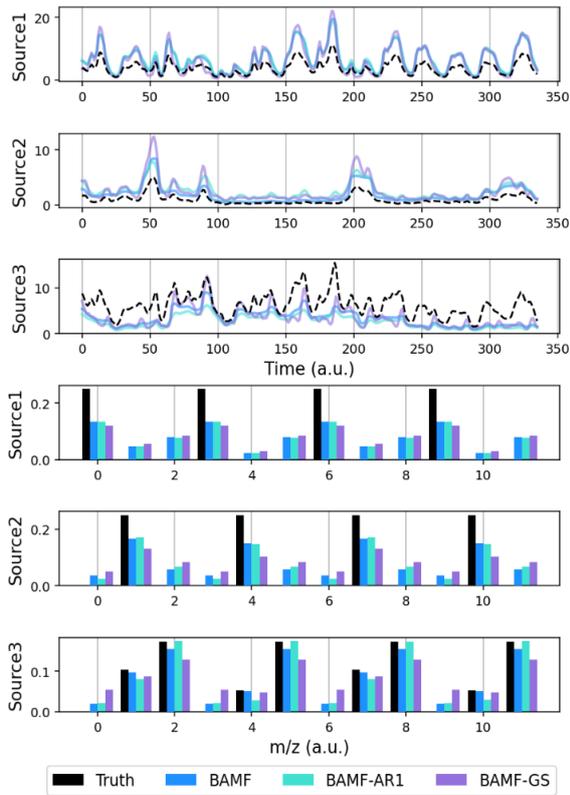
28



29

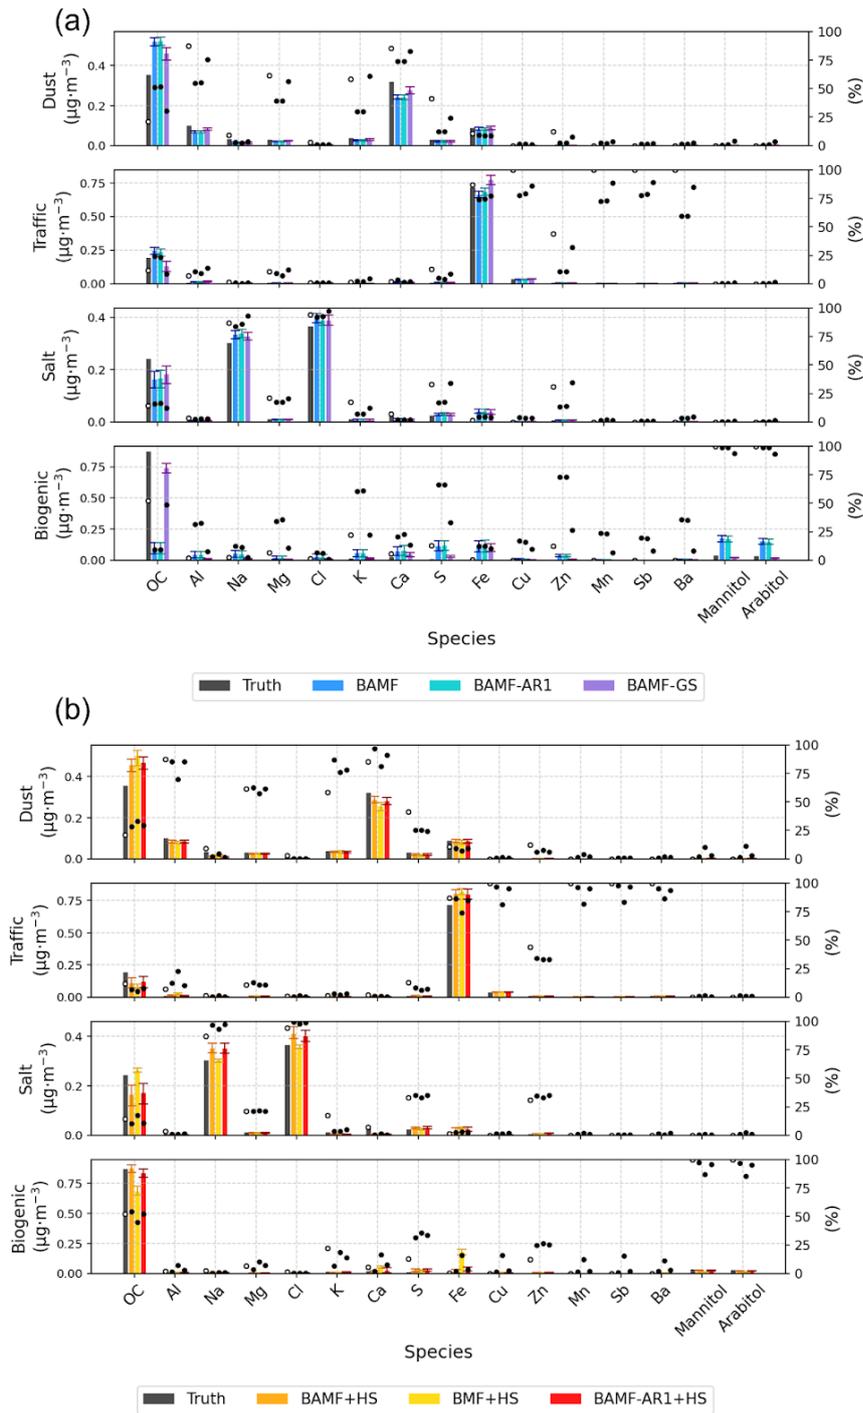
30 **Figure S6. Comparison of toy dataset distribution of F components for BAMF+HS, BAMF+Lasso, BAMF+Spike-and-**
 31 **slab.**

32



33

34 **Figure S7. Left: Toy dataset time series (top) and profiles (bottom) results with BAMF, BAMF-AR1 and BAMF-GS.**
 35 **Right: Toy dataset distribution of F components for BAMF, BAMF+HS, BAMF-AR1, BAMF-AR1+HS.**



36

37 **Figure S8. Synthetic offline dataset comparison between (a) BAMF, BAMF-AR1, BAMF-GS (b) BAMF+HS, BMF+HS,**
 38 **BAMF-AR1+HS.**

39

40 **Table S4. Purely-measurement-based offline synthetic dataset factorisation statistics. The last three rows correspond**
 41 **to the sum of the metrics for all all factors, except for the G/G₀ metric column, which shows the sum of the factors G/G₀**
 42 **deviation to 1.**

43

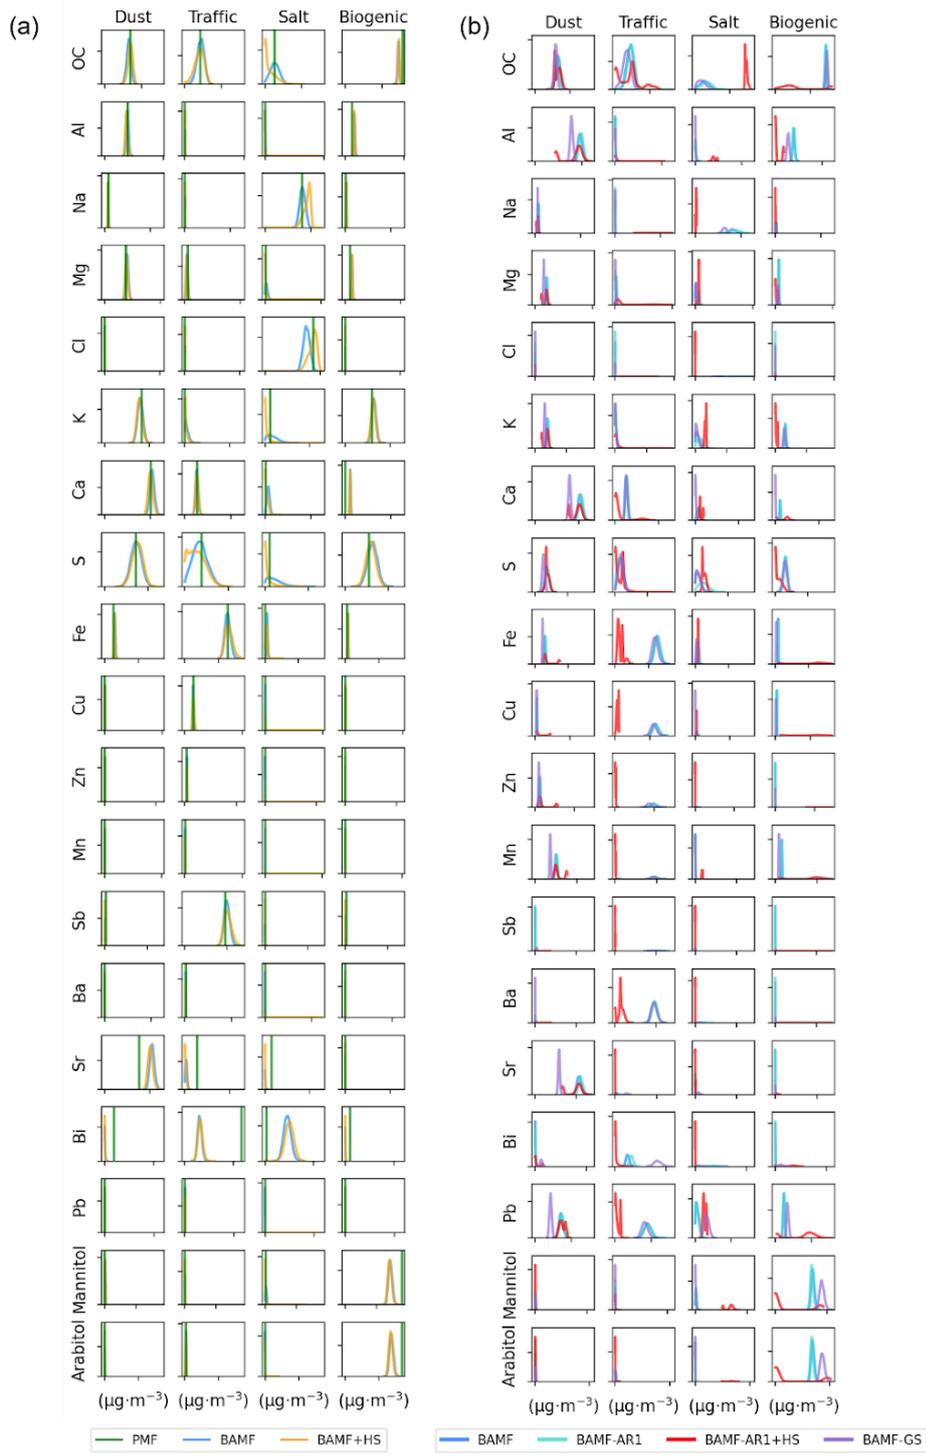
Model	Factor	F R ²	F ρ	F contr. R ²	F contr. ρ	F Gini	G R ²	G/G ₀ (*)
-------	--------	---------------------	--------	----------------------------	---------------	-----------	---------------------	-------------------------

Dust	PMF	0.98	0.90	0.93	0.90	0.77	0.95	1.00
	BAMF	0.99	0.97	0.78	0.93	0.74	0.92	0.98
	BAMF+HS	0.96	0.96	0.96	0.86	0.77	0.93	1.02
Traffic	PMF	0.99	0.67	0.87	0.88	0.89	0.94	0.83
	BAMF	0.97	0.77	0.46	0.75	0.81	0.93	1.02
	BAMF+HS	0.96	0.85	0.85	0.94	0.88	0.94	0.75
Salt	PMF	0.90	0.60	0.95	0.63	0.83	0.77	0.92
	BAMF	0.90	0.92	0.71	0.63	0.63	0.30	0.11
	BAMF+HS	0.78	0.82	0.92	0.55	0.86	0.52	0.86
Biogenic	PMF	0.96	0.27	0.77	0.16	0.88	0.96	0.68
	BAMF	0.03	0.85	0.56	0.21	0.51	0.96	0.14
	BAMF+HS	0.999	0.63	0.85	0.41	0.89	0.96	0.65
Σ_k (*)	PMF	3.85	2.44	3.52	2.57	3.37	3.64	0.61
	BAMF	2.89	3.51	2.51	2.52	2.69	3.11	1.79
	BAMF+HS	3.70	3.26	3.58	2.76	3.40	3.35	0.78

44

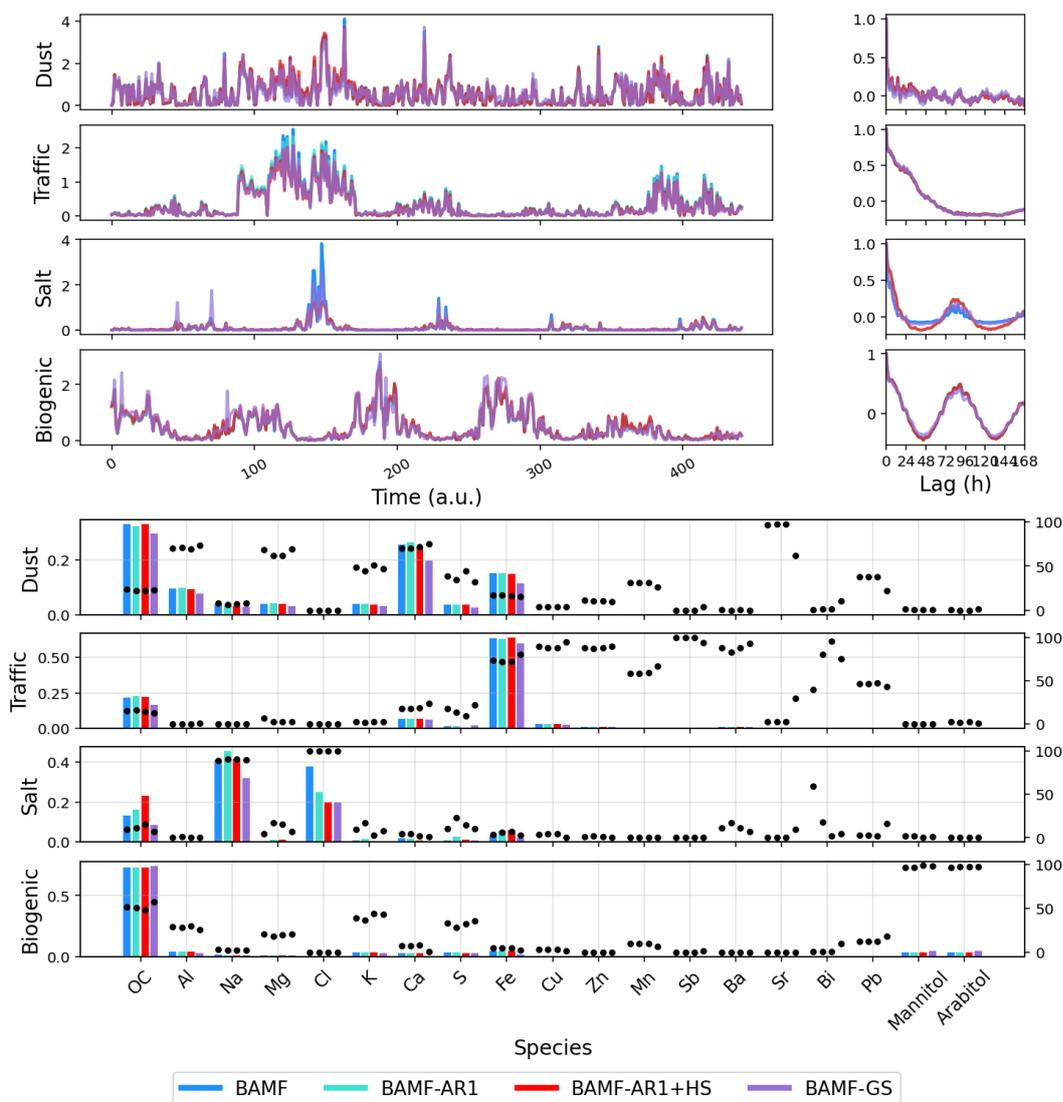
45

46



47
 48
 49
 50
 51
 52

Figure S9. Filters real-world filters dataset profiles components distributions for (a) PMF, BAMF, BAMF+HS; (b) BAMF, BAMF-AR1, BAMF-AR1+HS, BAMF-GS.



53

54 **Figure S10. Comparison of BAMF, BAMF-AR1, BAMF-AR1+HS, BAMF-GS on the real-world filters dataset. Plots**
 55 **from left to right and from the top to the bottom are: Time series, autocorrelation, profiles.**

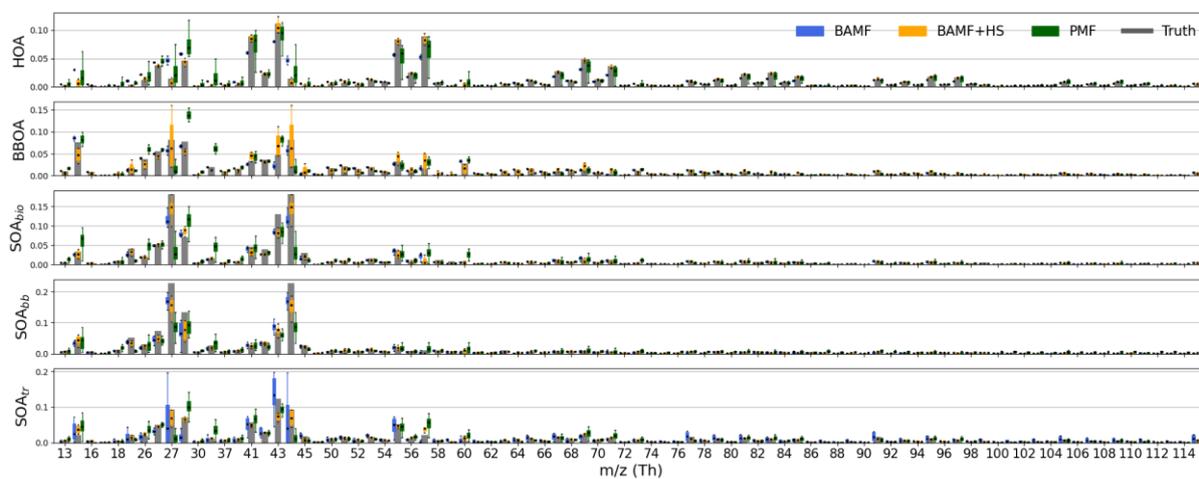
56

57 **Table S5. Online European synthetic datasets reconstruction and factorisation statistics. Statistics contain data from**
 58 **all sites, sources, and datasets.**

Statistic	Model	Mean	Median	Min	Max
$ X - Z /\sigma$	PMF	4.24	4.34	1.45	7.17
	BAMF	0.64	0.66	0.04	2.04
	BAMF+HS	2.92	2.26	0.66	8.45
G/G_0	PMF	2.82	1.10	0.39	18.45
	BAMF	1.07	0.92	0.02	7.21
	BAMF+HS	2.00	1.11	0.19	46.86

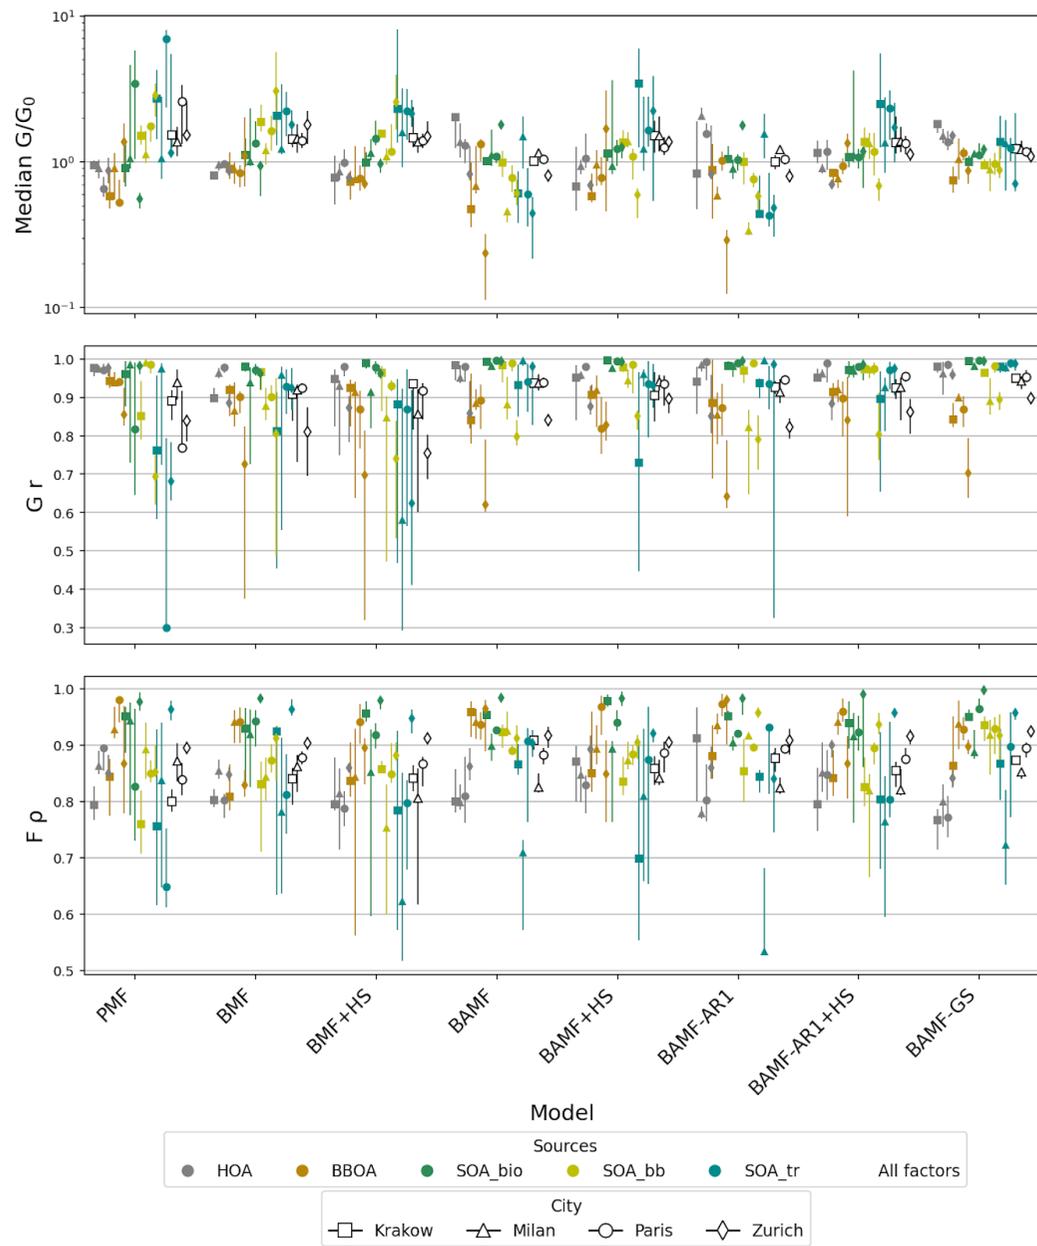
G R²	PMF	0.78	0.91	0.01	1.00
	BAMF	0.84	0.95	0.04	1.00
	BAMF+HS	0.83	0.92	0.05	1.00
F ρ	PMF	0.85	0.87	0.53	1.00
	BAMF	0.88	0.91	0.27	0.99
	BAMF+HS	0.87	0.90	0.53	1.00
F R²	PMF	0.84	0.92	0.04	1.00
	BAMF	0.90	0.96	0.02	1.00
	BAMF+HS	0.88	0.95	0.16	1.00
F sparsity	PMF	2.84	2.69	1.69	5.14
	BAMF	2.61	2.45	1.53	4.53
	BAMF+HS	2.63	2.45	1.65	4.53
F Gini	PMF	0.59	0.59	0.45	0.72
	BAMF	0.57	0.56	0.05	0.75
	BAMF+HS	0.57	0.58	0.41	0.83
F Gini ratio	PMF	0.99	0.96	0.75	1.46
	BAMF	0.97	0.99	0.09	1.15
	BAMF+HS	0.98	0.97	0.72	1.40

59



60

61 **Figure S11. Example results obtained with BAMF, BAMF+HS, and PMF for one of the 24 datasets included**
62 **in the ACSM-like experiment. The results correspond to Zurich, covering the period from 1 January 2019**
63 **to 14 January 2019 (dataset 0).**



64

65 **Figure S12. Summary metrics for all synthetic datasets.**

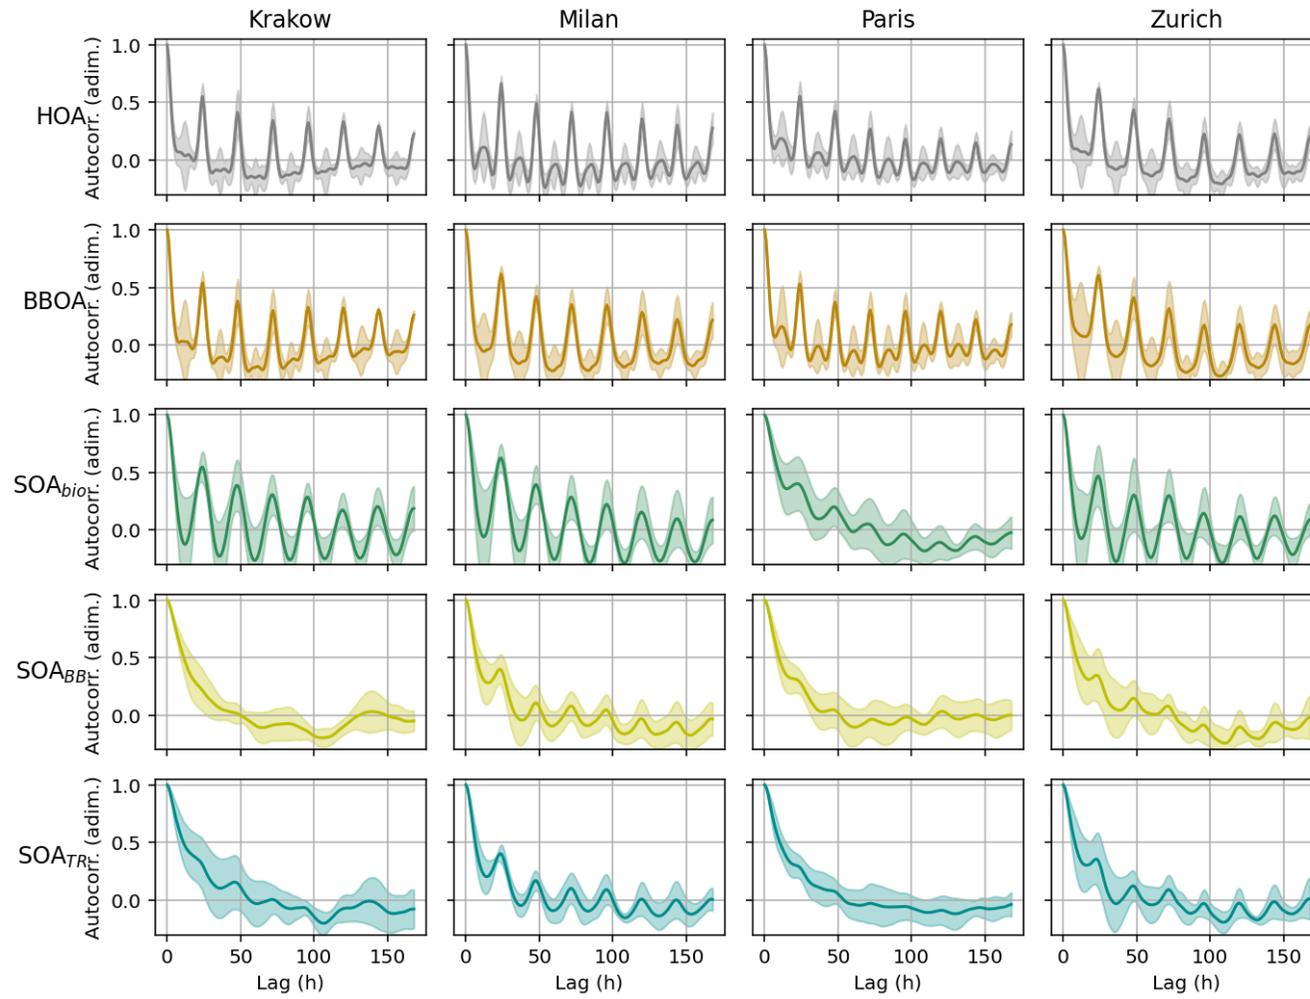
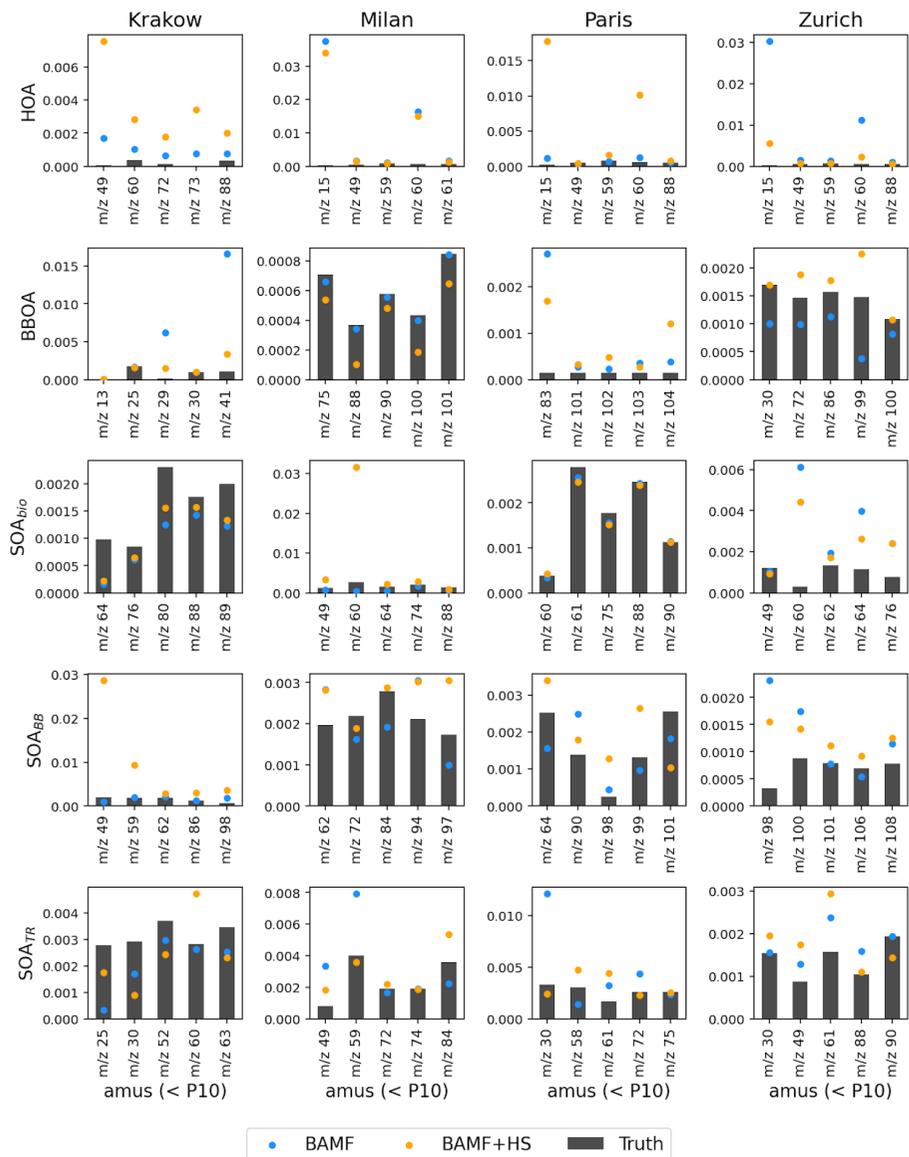


Figure S13. Autocorrelation means (solid lines) and standard deviations (shaded areas) of the 6 datasets per city and source.

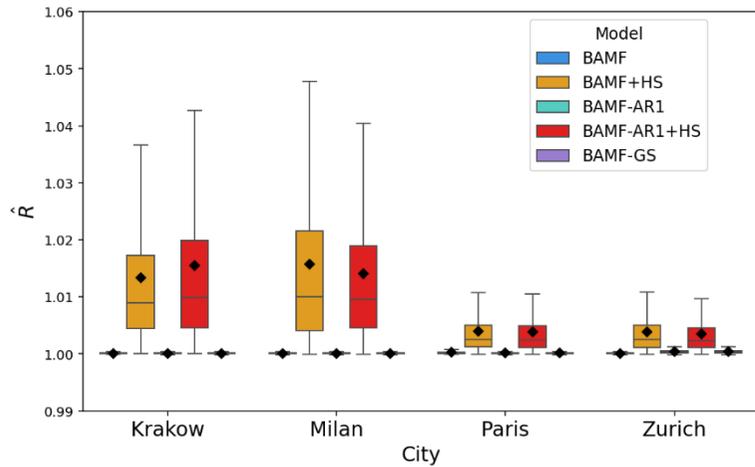


69

70

71

Figure S14. Sparsity introduction results for BAMF and BAMF+HS models in the European synthetic datasets for the 5 less massive m/zs, below the percentile 10. The bars show the truth concentration for the given m/z, and BAMF, BAMF+HS results are shown in markers.



72
73
74 **Figure S15. \hat{R} for the different cities and models. Model boxplots are in the same order as the legend.**

75
76 **2. Supplementary model formulation**

77 **Lasso**

78 The Lasso (Least Absolute Shrinkage and Selection Operator, Rasmussen et al. 2012) is a regularization method that adds an L1 penalty to linear regression, shrinking some coefficients
79 exactly to zero. The L1 penalty is a way to discourage large coefficients in a regression model by adding up the absolute values of all the model's coefficients and including that total in the
80 cost the model tries to minimize. This encourages sparsity in the model, making it useful for variable selection. However, it can over-shrink large coefficients and it treats all coefficients
81 equally, regardless of their signal strength, hence it tends to over-shrink even large signals.

82 **Spike-and-slab**

83 The spike-and-slab prior (Andersen et al. 2014) is a Bayesian approach to variable selection that models each coefficient as coming from a mixture of two distributions: a spike at zero
84 (forcing sparsity) and a slab (a wider distribution allowing nonzero values). It provides strong interpretability and explicit variable inclusion, but is computationally intensive due to its
85 discrete model space. This working scheme provides a binary shrinkage behaviour since the signals are considered either zero or slabs, which enforces 0-like signals, however, the selection
86 of non-zero elements can be too harsh.

87

88 **Horseshoe**

89 The regularized horseshoe prior is an extension of the standard horseshoe prior designed to improve robustness and regularization in sparse Bayesian models. Like the original horseshoe,
90 it combines global shrinkage (for overall sparsity) and local shrinkage (for individual coefficients), using heavy tails to allow large signals while aggressively shrinking noise. What sets
91 the regularized version apart is the addition of a slab component that limits the influence of extremely large coefficients. This slab acts like a soft ceiling, preventing the model from over-
92 trusting outlier variables while maintaining flexibility. This makes the prior more stable in practice, especially in finite data settings or when outliers are present. In comparison to the two
93 above priors, the regularized horseshoe provides a middle ground: it allows for adaptive shrinkage, handles strong signals well, and includes a slab to avoid the horseshoe's instability. It's
94 more robust and scalable than spike-and-slab and more flexible and statistically principled than Lasso.

96 **3. Supplementary model evaluation**

97 **3.1. Chemically sparse toy dataset**

98
99 Some alternative autocorrelation formulations to the regular BAMF autocorrelation priors were tried for the toy dataset. BMF, BMF+HS, and BMF-GS – analogs to BAMF, BAMF+HS,
100 BAMF-GS, respectively, but without the autocorrelation term – provided unstable results, with the highest $|\mathbf{Z}-\mathbf{X}|/\sigma$ median and maxima and the poorest \mathbf{X} vs. \mathbf{Z} correlation coefficients,
101 hence, these were discarded for this discussion. Figure S7 shows the time series, profiles and \mathbf{F} components distributions for BAMF, BAMF-AR1, BAMF-GS, BAMF+HS. Both regarding
102 time series and profiles, BAMF seems to be resembling the truth equally or better than BAMF-AR1 and BAMF-GS with the only exception of the time series of Source 3, which is visually
103 better captured by BAMF-GS. However, BAMF-GS profiles are further from the truth and m/z mass closure for this model is much worse than regular BAMF, as also seen by poorer
104 Spearman and Pearson correlation coefficients. Profiles as captured by BAMF-AR1 are not consistently better than those from BAMF either. BAMF appears as the most balanced and
105 accurate model of the three for the presented toy dataset. Regarding \mathbf{F} components distributions, BAMF and BAMF-AR1 and BAMF+HS and BAMF-AR1+HS BAMF-AR1, BAMF-
106 AR1+HS are compared in the right panel of Figure S7. BAMF and BAMF-AR1 show very similar distributions, with the slight differences mentioned before. BAMF+HS and BAMF-
107 AR1+HS present some differences, and even if the shrinkage is also present in BAMF-AR1+HS (although with weaker Gini than BAMF+HS), some components present distributions with
108 a more pronounced multimodal behaviour than the BAMF+HS. Therefore, the BAMF-AR1+HS is shown to be less precise and accurate than BAMF+HS. The BAMF-GS+HS was not tried
109 out since the high concentrations of the components of the non-normalised \mathbf{F} matrix on this model hinder their shrinkage to zero.

111 **3.2 Chemically sparse offline synthetic dataset**

112

113 Figure S8 shows the profiles of the \mathbf{F} matrix for BAMF, BAMF-AR1, BAMF-GS (a) and BAMF+HS, BMF+HS, BAMF-AR1+HS (b) models. BAMF, BAMF-AR1, and BAMF-GS were
114 run, initialising \mathbf{F} as a normal distribution to make the sampling more sturdy to avoid initialization failure. The outcomes of these models show general good agreement with the truth, with
115 BAMF+HS being the most accurate model followed by BAMF-AR1+HS. The sum across factors of profile Pearson R^2 with truth for BAMF+HS, BMF+HS, BAMF-AR1+HS, and BAMF-
116 GS were 3.90, 3.80, and 3.90, proving the beneficial effect of the horseshoe prior in the description of sparse profiles. The BAMF-AR1+HS model is performing slightly worse than the
117 BAMF+HS in time series (G/G_0 deviation sum of 0.85, 0.91, respectively). Regarding the non-horseshoed models, the sum of the correlation with truth \mathbf{G} for all factors was very similar
118 for BAMF, BAMF-AR1, BAMF-GS (3.93, 3.92, 3.93, respectively) but the Spearman correlation coefficient with truth \mathbf{F} was better for BAMF (3.75, 3.71, 3.67, respectively), highlighting
119 a slight better performance for BAMF. Hence, with all, the BAMF+HS seems the most accurate model, benefitting both from its autocorrelation and sparsity properties.

120

121 3.3 Chemically sparse offline real-world dataset

122 Figure S10 shows the performance of the other models discussed in the Toy dataset and European city datasets sections, BAMF-AR1, BAMF-AR1+HS, BAMF-GS, in comparison to
123 BAMF, used as the base case. All models employed seem to agree, providing overlapping time series and autocorrelations, and similar profiles, with the exception of BAMF-AR1+HS,
124 which is slightly differentiated from the others. However, it does not provide sparsified profiles in the species in which the shrinkage effect was found for BAMF+HS or elsewhere. This is
125 more evidently depicted in Figure S9 (b), in which the distributions of \mathbf{F} components are shown. The BAMF-AR1+HS does not present the expected sparsifying shape except for Ca, K,
126 and some elements in the traffic and salt profiles. In general, the BAMF-AR1+HS distributions are more multimodal, reflecting higher divergence across HMC chains which makes results
127 much less sturdy. Regarding the rest of models, the distributions are very similar ensuring the robustness of all three and stability of the solution. Therefore, all Bayesian models and PMF
128 point to a similar, robust solution for the filters dataset. The horseshoe prior addition to BAMF, though, provides here a useful sparsity introduction in the current dataset which helps purify
129 the profiles from unwanted profile entanglement. However, the implementation of this prior in the BAMF-AR1 model detracts the solution due to HMC chain divergence.

130

131 References

132 Crippa, M., DeCarlo, P. F., Slowik, J. G., Mohr, C., Heringa, M. F., Chirico, R., Poulain, L., Freutel, F., Sciare, J., Cozic, J., Di Marco, C. F., Elsasser, M., Nicolas, J. B., Marchand, N.,
133 Abidi, E., Wiedensohler, A., Drewnick, F., Schneider, J., Borrmann, S., Nemitz, E., Zimmermann, R., Jaffrezo, J.-L., Prévôt, A. S. H., and Baltensperger, U.: Wintertime aerosol chemical
134 composition and source apportionment of the organic fraction in the metropolitan area of Paris, *Atmos. Chem. Phys.*, 13, 961–981, <https://doi.org/10.5194/acp-13-961-2013>, 2013.
135 Daellenbach, K. R., Manousakas, M., Jiang, J., Cui, T., Chen, Y., El Haddad, I., Fermo, P., Colombi, C., and Prévôt, A. S. H.: Organic aerosol sources in the Milan metropolitan area–
136 Receptor modelling based on field observations and air quality modelling. *Atmospheric Environment*, 307, 119799, <https://doi.org/10.1016/j.atmosenv.2023.119799>, 2023.
137 Elser, M., Bozzetti, C., El-Haddad, I., Maasikmets, M., Teinmaa, E., Richter, R., Wolf, R., Slowik, J. G., Baltensperger, U., and Prévôt, A. S. H.: Urban increments of gaseous and aerosol
138 pollutants and their sources using mobile aerosol mass spectrometry measurements, *Atmos. Chem. Phys.*, 16, 7117–7134, <https://doi.org/10.5194/acp-16-7117-2016>, 2016.

139 Mohr, C., DeCarlo, P. F., Heringa, M. F., Chirico, R., Slowik, J. G., Richter, R., Reche, C., Alastuey, A., Querol, X., Seco, R., Peñuelas, J., Jiménez, J. L., Crippa, M., Zimmermann, R.,
140 Baltensperger, U., and Prévôt, A. S. H.: Identification and quantification of organic aerosol from cooking and other sources in Barcelona using aerosol mass spectrometer data, *Atmos.*
141 *Chem. Phys.*, 12, 1649–1665, <https://doi.org/10.5194/acp-12-1649-2012>, 2012.

142 Tobler, A. K., Skiba, A., Canonaco, F., Močnik, G., Rai, P., Chen, G., Bartyzel, J., Zimnoch, M., Styszko, K., Nęcki, J., Furger, M., Róžański, K., Baltensperger, U., Slowik, J. G., and
143 Prevot, A. S. H.: Characterization of non-refractory (NR) PM₁ and source apportionment of organic aerosol in Kraków, Poland, *Atmos. Chem. Phys.*, 21, 14893–14906,
144 <https://doi.org/10.5194/acp-21-14893-2021>, 2021.

145 Ulbrich, I. M., Canagaratna, M. R., Zhang, Q., Worsnop, D. R., and Jimenez, J. L.: Interpretation of organic components from Positive Matrix Factorization of aerosol mass spectrometric
146 data, *Atmos. Chem. Phys.*, 9, 2891–2918, <https://doi.org/10.5194/acp-9-2891-2009>, 2009.

147 Ulbrich, I. M., Handschy, A., Lechner, M., and Jimenez, J. L.: AMS Spectral Database, <http://cires.colorado.edu/jimenez-group/AMSsd/>, last access: 19 May 2025.

148 Via, M., Minguillón, M. C., Reche, C., Querol, X., and Alastuey, A.: Increase in secondary organic aerosol in an urban environment, *Atmos. Chem. Phys.*, 21, 8323–8339,
149 <https://doi.org/10.5194/acp-21-8323-2021>, 2021.

150 Zhang, Y., Favez, O., Petit, J.-E., Canonaco, F., Truong, F., Bonnaire, N., Crenn, V., Amodeo, T., Prévôt, A. S. H., Sciare, J., Gros, V., and Albinet, A.: Six-year source apportionment of
151 submicron organic aerosols from near-continuous highly time-resolved measurements at SIRTa (Paris area, France), *Atmos. Chem. Phys.*, 19, 14755–14776, [https://doi.org/10.5194/acp-](https://doi.org/10.5194/acp-19-14755-2019)
152 [19-14755-2019](https://doi.org/10.5194/acp-19-14755-2019), 2019.