

# Rational Ignorance: Simpler Models Learn More Information from Finite Data

Henry H. Mattingly,<sup>1</sup> Mark K. Transtrum,<sup>2</sup> Michael C. Abbott,<sup>3,\*</sup> and Benjamin B. Machta<sup>4,†</sup>

<sup>1</sup>*Lewis-Sigler Institute and Department of Chemical and Biological Engineering,  
Princeton University, Princeton, NJ 08544, USA.*

<sup>2</sup>*Department of Physics and Astronomy, Brigham Young University, Provo, Utah 84602, USA.*

<sup>3</sup>*Institute of Physics, Jagiellonian University, Ulica Łojasiewicza 11, 30-348 Kraków, Poland.*

<sup>4</sup>*Lewis-Sigler Institute and Department of Physics,  
Princeton University, Princeton, NJ 08544, USA.*

(Dated: 2 May 2017, updated 31 August 2017)

We use the language of uninformative Bayesian prior choice to study the selection of appropriately simple effective models. We advocate for the prior which maximizes the mutual information between parameters and predictions, learning as much as possible from limited data. When many parameters are poorly constrained by the available data, we find that this prior puts weight only on boundaries of the parameter manifold. Thus it selects a lower-dimensional effective theory in a principled way, ignoring irrelevant parameter directions. In the limit where there is sufficient data to tightly constrain any number of parameters, this reduces to Jeffreys prior. But we argue that this limit is pathological, as it leads to dramatic dependence on effects invisible to experiment.

Physicists prefer simple models not because nature is simple, but because most of its complication is usually irrelevant. Our most rigorous understanding of this idea comes from the Wilsonian renormalization group [1–3], which describes mathematically the process of zooming out and losing sight of microscopic details. These details only influence the effective theory which describes macroscopic observables through a few relevant parameter combinations, such as the critical temperature, or the proton mass. The remaining irrelevant parameters can be ignored, as they are neither constrained by past data nor useful for predictions. Such models can now be understood as part of a large class called sloppy models [4–13], whose usefulness relies on a similar compression of a large microscopic parameter space down to just a few relevant directions.

This justification for model simplicity is different from the one more often discussed in statistics, motivated by the desire to avoid overfitting [14–20]. Since irrelevant parameters have an almost invisible effect on predicted data, they cannot be excluded on these grounds. Here we motivate their exclusion differently: we show that simplifying a model can often allow it to extract more information from a limited data set, and that this offers a guide for choosing appropriate effective theories.

We phrase the question of model selection as part of the choice of a Bayesian prior on some high-dimensional parameter space. In a set of nested models, we can always move to a simpler model by using a prior which is nonzero only on some subspace. Recent work has suggested that interpretable effective models are typically obtained by taking some parameters to their limiting values, often 0 or  $\infty$ , thus restricting to lower-dimensional boundaries of

the parameter manifold [21].

Our setup is that we wish to learn about a theory by performing some experiment which produces data  $x \in X$ . The theory and the experiment are together described by a probability distribution  $p(x|\theta)$ , for each value of the theory’s parameters  $\theta \in \Theta$ . This function encodes both the quality and quantity of data to be collected.

The mutual information between the parameters and their expected data is defined as  $MI = I(X; \Theta) = S(\Theta) - S(\Theta|X)$ , where  $S$  is the Shannon entropy [22]. The MI thus quantifies the information which can be learned about the parameters by measuring the data, or equivalently, the information about the data which can be encoded in the parameters [23, 24]. Defining  $p_*(\theta)$  by maximizing this information, we find:

- i. The prior  $p_*(\theta)$  is almost always discrete, with weight only on a finite number  $K$  of points, or atoms (Figure 1):  $p_*(\theta) = \sum_{a=1}^K \lambda_a \delta(\theta - \theta_a)$ .
- ii. When data is abundant,  $p_*(\theta)$  approaches Jeffreys prior  $p_J(\theta)$ . As this continuum limit is approached, the proper spacing of the atoms shrinks as a power law (Figure 2).
- iii. When data is scarce, most atoms lie on boundaries of parameter space, corresponding to effective models with fewer parameters (Figure 3). The resulting distribution of weight along relevant directions is much more even than that given by Jeffreys prior (Figure 4).

After some preliminaries, we demonstrate these properties in three simple examples, stylized versions of realistic experiments. To see the origin of discreteness, we study the bias of an unfair coin and then a single variable corrupted with Gaussian noise. To see how models of lower dimension arise, we then study the problem of inferring decay rates in a sum of exponentials.

\* [abbott@th.if.uj.edu.pl](mailto:abbott@th.if.uj.edu.pl)

† [bmachta@princeton.edu](mailto:bmachta@princeton.edu)

## Priors and Geometry

Bayes' theorem tells us how to update our knowledge of  $\theta$  upon observing data  $x$ , from prior  $p(\theta)$  to posterior  $p(\theta|x) = p(x|\theta)p(\theta)/p(x)$ , where  $p(x) = \int d\theta p(\theta)p(x|\theta)$ . In the absence of better knowledge we must pick an uninformative prior which codifies our ignorance. The naive choice of a flat prior  $p(\theta) = \text{const.}$  has undesirable features, in particular making  $p(x)$  depend on the choice of parameterization, through the measure  $d\theta$ .

Demanding reparameterization invariance led Jeffreys to the prior  $p_J(\theta) \propto \sqrt{\det g_{\mu\nu}}$  [25] which is, up to normalization, the volume form arising from the Fisher information metric (FIM, often -matrix):

$$g_{\mu\nu}(\vec{\theta}) = \int dx p(x|\vec{\theta}) \frac{\partial \log p(x|\vec{\theta})}{\partial \theta^\mu} \frac{\partial \log p(x|\vec{\theta})}{\partial \theta^\nu}.$$

This Riemannian metric defines a reparameterization-invariant distance between points,  $ds^2 = \sum_{\mu,\nu=1}^D g_{\mu\nu} d\theta^\mu d\theta^\nu$ . It measures the distinguishability of the data they are expected to produce, in units of standard deviations. Repeating an (identical and independently distributed) experiment  $m$  times means considering  $p^m(\vec{x}|\theta) = \prod_{j=1}^m p(x_j|\theta)$ , which leads to metric  $g_{\mu\nu}^m(\theta) = m g_{\mu\nu}(\theta)$ . However the factor  $m^{D/2}$  in the volume is lost by normalizing  $p_J(\theta)$ . Thus Jeffreys prior depends on the type of experiment, but not the quantity of data.

Bernardo defined a prior  $p_\star(\theta)$  by maximizing the mutual information between parameters  $\Theta$  and the expected data  $X^m$  from  $m$  repetitions, and then a reference prior by taking the limit  $m \rightarrow \infty$  [26, 27]. Under certain benign assumptions, this reference prior is exactly Jeffreys prior [26, 28, 29], providing an alternative justification for  $p_J(\theta)$ .

We differ in taking seriously that the amount of data collected is always finite.<sup>1</sup> Besides being physically unrealistic, the limit  $m \rightarrow \infty$  is pathological both for model selection and prior choice. In this limit any number of parameters can be perfectly inferred, justifying an arbitrarily complicated model. In addition, in this limit the posterior  $p(\theta|x)$  becomes independent of any smooth prior.<sup>2</sup>

Geometrically, the defining feature of sloppy models is that they have a parameter manifold with hyper-ribbon structure [6–9]: there are some long directions (corresponding to  $d$  relevant, or stiff, parameters) and many shorter directions ( $D - d$  irrelevant, or sloppy, parameter combinations). These lengths are often estimated using

the eigenvalues of  $g_{\mu\nu}$ , and have logarithms about evenly-spaced over many orders of magnitude [4, 5]. The effect of coarse-graining is to shrink irrelevant directions (here using the technical meaning of irrelevant: a parameter which shrinks under renormalization group flow) while leaving relevant directions extended, producing a sloppy manifold [8]. By contrast the limit  $m \rightarrow \infty$  has the effect of expanding all directions, thus erasing the distinction between directions longer and shorter than the critical length scale of (approximately) one standard deviation.

On such a hyper-ribbon, Jeffreys prior has an undesirable feature: since it is constructed from the  $D$ -dimensional notion of volume, its weight along the relevant directions always depends on the volume of the  $D - d$  irrelevant directions. This gives it extreme dependence on which irrelevant parameters are included in the model.<sup>3</sup> The optimal prior  $p_\star(\theta)$  avoids this dependence because it is almost always discrete, at finite  $m$ .<sup>4</sup> It puts weight on a set of nearly distinguishable points, closely spaced along the relevant directions, but ignoring the irrelevant ones. Yet being the solution to a reparameterization-invariant optimization problem, the prior  $p_\star(\theta)$  retains this good feature of  $p_J(\theta)$ .<sup>5</sup>

Maximizing the mutual information was originally done to calculate the capacity of a communication channel, and we can borrow techniques from rate-distortion theory here. In particular, the discreteness we exploit was discovered several times in engineering [33] [34] [35, 36] before being discussed in statistics [27, 29, 37, 38]. We adapt our title from the idea of rational inattention in economics, wherein market actors have a finite bandwidth for news, driving them to make discrete choices [39, 40]. Rate-distortion theory has also been useful in several areas of biology [41–43], and discreteness emerges in a recent theoretical model of the immune system [44].

<sup>3</sup> See Figure 4 for a demonstration of this point. For another example, consider a parameter manifold  $\Theta$  which is a cone, with Fisher metric  $ds^2 = (50 d\vartheta)^2 + \vartheta^2 d\Omega_n^2/4$ : there is one relevant direction  $\vartheta \in [0, 1]$  of length  $L = 50$ , and  $n$  irrelevant directions forming a sphere of diameter  $\vartheta$ . Then the prior on  $\vartheta$  alone implied by  $p_J(\theta)$  is  $p(\vartheta) = (n+1)\vartheta^n$ , putting most of the weight near to  $\vartheta = 1$ , dramatically so if  $n = D - d$  is large. But since only the relevant direction is visible to our experiment, the region  $\vartheta \approx 0$  ought to be treated similarly to  $\vartheta \approx 1$ . Our prior  $p_\star(\theta)$  has this property.

<sup>4</sup> We offer both numerical and analytic arguments for discreteness below. The exception to discreteness is that if there is an exact continuous symmetry,  $p_\star(\theta)$  will be constant along it. For example if our Gaussian model (2) is placed on a circle (identifying both  $\theta \sim \theta + 1$  and  $x \sim x + 1$ ) then the optimum prior is a constant.

<sup>5</sup> There are some similarities to the penalised complexity priors of [32]. Those are also reparameterization invariant, and also concentrate weight on a subspace of  $\Theta$ , often a boundary. However both the subspace (or base model) and the degree of concentration (scaling parameter) are chosen by hand, rather than being deduced from  $p(x|\theta)$ .

<sup>1</sup> Interned for five years, John Kerrich only flipped his coin  $10^4$  times [30]. With computers we can do better, but even the LHC only generated about  $10^{18}$  bits of data [31].

<sup>2</sup> For simplicity we consider only regular models, i.e. we assume all parameters are structurally identifiable.




Figure 1. **Optimal priors for the Bernoulli model (1).** Red lines indicate the positions of delta functions in  $p_*(\theta)$ , which are at the maxima of  $f_{\text{KL}}(\theta)$ , sometimes called the Bayes risk. As  $m \rightarrow \infty$  these coalesce into Jeffreys prior  $p_J(\theta)$ .

### One-Parameter Models

The first example we study is the Bernoulli problem, in which we wish to determine the probability  $\theta \in [0, 1]$  that an unfair coin gives heads, using the data from  $m$  trials. It is sufficient to record the total number of heads  $x$ , which occurs with probability

$$p(x|\theta) = \frac{m!}{x!(m-x)!} \theta^x (1-\theta)^{m-x}. \quad (1)$$

This gives  $g_{\theta\theta} = \frac{m}{\theta(1-\theta)}$ , thus  $p_J(\theta) = [\pi\sqrt{\theta(1-\theta)}]^{-1}$ , and proper parameter space length  $L = \int \sqrt{ds^2} = \pi\sqrt{m}$ .

In the extreme case  $m = 1$ , the optimal prior is two delta functions,  $p_*(\theta) = \frac{1}{2}\delta(\theta) + \frac{1}{2}\delta(\theta-1)$ , and  $\text{MI} = \log 2$ , exactly one bit [27, 29, 37]. Before an experiment that will only ever be run once, this places equal weight on both outcomes; afterwards it records the outcome. As  $m$  increases, weight is moved from the boundary onto interior points, which increase in number, and ultimately approach the smooth  $p_J(\theta)$ : see Figures 1 and 2A.

Similar behavior is seen in a second example, in which we measure one real number  $x$ , normally distributed with known  $\sigma$  about the parameter  $\theta \in [0, 1]$ :

$$p(x|\theta) = \frac{1}{\sqrt{2\pi}\sigma} e^{-(x-\theta)^2/2\sigma^2}. \quad (2)$$

Repeated measurements are equivalent to smaller  $\sigma$  (by  $\sigma \rightarrow \sigma/\sqrt{m}$ ), so we fix  $m = 1$  here. The Fisher metric is  $g_{\theta\theta} = 1/\sigma^2$ , thus  $L = 1/\sigma$ . An optimal prior is shown in Figure 4A, along with its implied distribution of expected data. This  $p(x)$  is similar to that implied by Jeffreys prior, here  $p_J(\theta) = 1$ .

We calculated  $p_*(\theta)$  numerically in two ways. After discretizing both  $\theta$  and  $x$ , we can use the Blahut–Arimoto (BA) algorithm [45, 46]. This converges to the global maximum, which is a discrete distribution: see Figure 5.

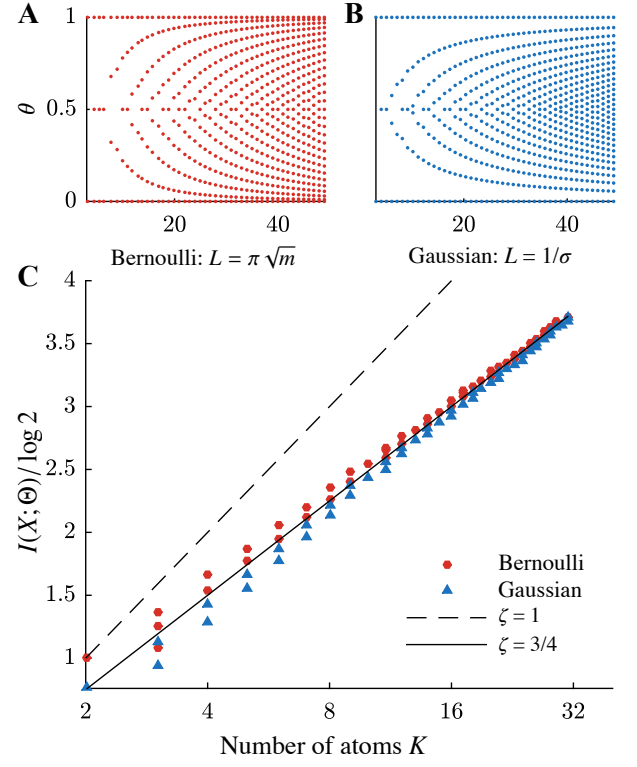


Figure 2. **Behavior of  $p_*(\theta)$  with increasing Fisher length.** Panels A and B show the atoms of  $p_*(\theta)$  for the two one-dimensional models as  $L$  is increased (i.e. we perform more repetitions  $m$  or have smaller noise  $\sigma$ ). Panel C shows the scaling of the mutual information (in bits) with the number of atoms  $K$ . The dashed line is the bound  $\text{MI} \leq \log K$ , and the solid line is the scaling law  $\text{MI} \sim 3/4 \log K$  leading to (4).

Alternatively, using our knowledge that  $p_*(\theta)$  is discrete, we can instead adjust the positions  $\theta_a$  and weights  $\lambda_a$  of a finite number of atoms. See the Appendix for more details.

To see analytically why discreteness arises, we write the mutual information as

$$\text{MI} = I(X; \Theta) = \int d\theta p(\theta) f_{\text{KL}}(\theta), \quad (3)$$

$$f_{\text{KL}}(\theta) = D_{\text{KL}}(p(x|\theta) \| p(x)) = \int dx p(x|\theta) \log \frac{p(x|\theta)}{p(x)}$$

where  $D_{\text{KL}}$  is the Kullback–Leibler divergence.<sup>6</sup> Maximizing MI over all functions  $p(\theta)$  with  $\int d\theta p(\theta) = 1$  gives  $f_{\text{KL}}(\theta) = \text{const}$ . But the maximizing function will not, in general, obey  $p(\theta) \geq 0$ . Subject to this inequality  $p_*(\theta)$  must satisfy

$$\{p_*(\theta) > 0, f_{\text{KL}}(\theta) = \text{MI}\} \text{ or } \{p_*(\theta) = 0, f_{\text{KL}}(\theta) < \text{MI}\}$$

<sup>6</sup> The function  $f_{\text{KL}}(\theta)$  is sometimes called the Bayes risk, as it quantifies how poorly the prior will perform if  $\theta$  turns out to be correct.

at every  $\theta$ . With finite data  $f_{\text{KL}}(\theta)$  – MI must be an analytic function of  $\theta$ , and therefore must be smooth with a finite numbers of zeros, corresponding to the atoms of  $p_*(\theta)$  (Figure 1A). See [27, 35, 40] for related arguments for discreteness.

The number of atoms occurring in  $p_*(\theta)$  increases as the data improves. For  $K$  atoms there is an absolute bound  $\text{MI} \leq \log K$ , saturated if they are perfectly distinguishable. In Figure 2C we observe that the optimal priors instead approach a line  $\text{MI} \rightarrow \zeta \log K$ , with slope  $\zeta < 1$ . At large  $L$  the length of parameter space is proportional to the number of distinguishable points, hence  $\text{MI} \rightarrow \log L$ . Together these imply  $K \sim L^{1/\zeta}$ , and so the average number density of atoms grows as

$$\rho_0 = K/L \sim L^{1/\zeta - 1} \approx L^{1/3}, \quad L \gg 1. \quad (4)$$

Thus the proper spacing between atoms shrinks to zero in the limit of infinite data, i.e. neighboring atoms cease to be distinguishable.

To derive this scaling law analytically, in a forthcoming paper we consider a field theory for the number density of atoms, in which the entropy density (omitting numerical factors) is

$$\mathcal{S} = \text{const.} - e^{-\rho^2} [\rho^4 (\rho')^2 + 1].$$

From this we find  $\zeta = 3/4$ , which is consistent with both models presented above.

### Multi-parameter Models

In the examples above,  $p_*(\theta)$  concentrates weight on the edges of its allowed domain when data is scarce (i.e. when  $m$  is small or  $\sigma$  is large, hence  $L$  is small). We next turn to a multi-parameter model in which some parameter combinations are ill-constrained, and where edges correspond to reduced models.

The physical picture is that we wish to determine the composition of an unknown radioactive source, from data of  $x_t$  Geiger counter clicks at some times  $t$ . As parameters we have the quantities  $A_\mu$  and decay constants  $k_\mu$  of isotopes  $\mu$ . The probability of observing  $x_t$  should be a Poisson distribution (of mean  $y_t$ ) at each time, but we approximate these by Gaussians to write:<sup>7</sup>

$$p(\vec{x}|\vec{y}) \propto \prod_t e^{-(x_t - y_t)^2 / 2\sigma_t^2}, \quad y_t = \sum_\mu A_\mu e^{-k_\mu t}. \quad (5)$$

Fixing  $\sigma_t = \sigma = \text{const.}$  then brings us to a nonlinear least-squares model of the type studied by [6, 7]. This same model also arises in other contexts, such as the

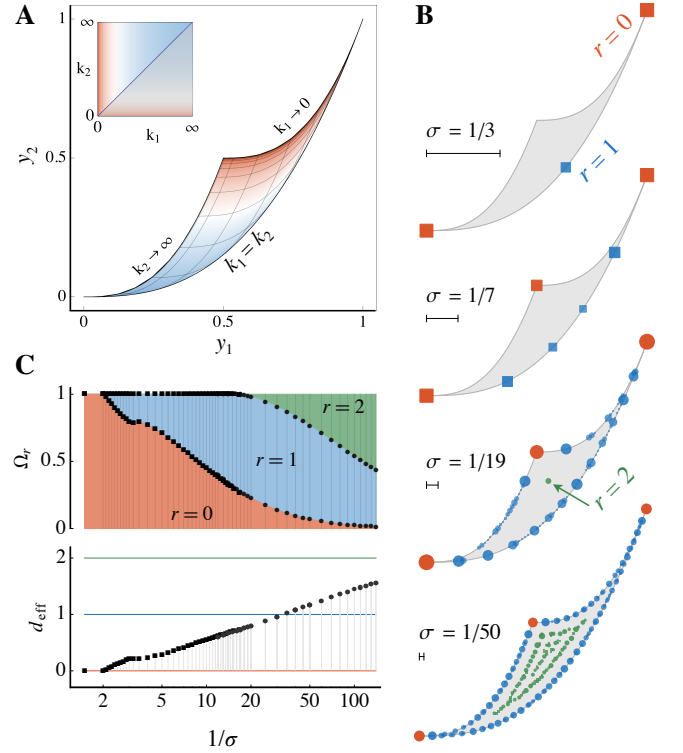


Figure 3. **Parameters and priors for the exponential model (5).** Panel A shows the area of the  $\vec{y}$  plane covered by all decay constants  $k_1, k_2 \geq 0$ . Notice that the initial model is slightly singular, since  $k_1 \leftrightarrow k_2$  does not change the predictions; we could cure this by restricting to  $k_2 \geq k_1$ , or by working with  $\vec{y}$ , to obtain a regular model. Panel B shows the positions of the delta functions of the optimal prior  $p_*(\vec{y})$  for several values of  $\sigma$ , with colors indicating the dimensionality  $r$  at each point. Panel C shows the proportion of weight on these dimensionalities.

asymptotic approach to equilibrium of many dynamical systems [47].

We can see the essential behavior with just two isotopes in fixed quantities:  $A_\mu = \frac{1}{2}$ , thus  $y_t = \frac{1}{2}(e^{-k_1 t} + e^{-k_2 t})$ . Measuring at only two times  $t_1$  and  $t_2$ , we almost have a two-dimensional version of (2), in which the center of the distribution  $\vec{y} = (y_1, y_2)$  plays the role of  $\theta$  above. The mapping between  $\vec{k}$  and  $\vec{y}$  is shown in Figure 3A, fixing  $t_2/t_1 = e$ . The Fisher information metric is proportional to the ordinary Euclidean metric for  $\vec{y}$ , but not for  $\vec{k}$ :

$$g_{\mu\nu}(\vec{k}) = \frac{1}{\sigma^2} \sum_t \frac{\partial y_t}{\partial k_\mu} \frac{\partial y_t}{\partial k_\nu} \iff g_{st}(\vec{y}) = \frac{1}{\sigma^2} \delta_{st}. \quad (6)$$

Thus Jeffreys prior is a constant on the allowed region of the  $\vec{y}$  plane.

Then we proceed to find the optimum  $p_*(\vec{y})$  for this model, shown in Figure 3B for various values of  $\sigma$ . When  $\sigma$  is large, this has delta functions only in two of the corners, allowing only  $k_1, k_2 = 0$  and  $k_1, k_2 = \infty$ . As  $\sigma$  is decreased, new atoms appear first along the lower boundary (corresponding to the one-dimensional model

<sup>7</sup> Using a normal distribution of fixed  $\sigma$  here is what allows the metric in (6) to be so simple. However the qualitative behavior from the Poisson distribution is very similar.



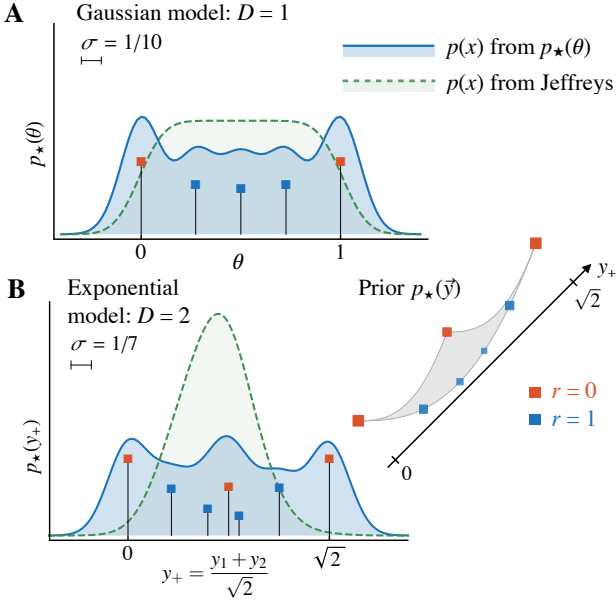


Figure 4. **Distributions of expected data  $p(x)$  from different priors.** Panel A is the one-parameter Gaussian model, with  $L = 10$ . Panel B projects the two-parameter exponential model onto the  $y_1 + y_2$  direction, for  $\sigma = 1/7$  where the perpendicular direction should be irrelevant. The length of the relevant direction is about the same as the one-parameter case:  $L_+ = 7\sqrt{2}$ . Notice that the distribution of expected data  $p(x_+)$  from Jeffreys prior here is quite different, with almost no weight at the ends of the range  $(0 \text{ and } \sqrt{2})$ , because this prior still weights the area not the length.

where  $k_1 = k_2$ ) and then along the other boundaries. At sufficiently small  $\sigma$ , atoms start filling in the (two-dimensional) interior.

To show this progression in Figure 3C, we define  $\Omega_r$  as the total weight on all edges of dimension  $r$ , and an effective dimensionality  $d_{\text{eff}} = \sum_{r=1}^D r \Omega_r$ . This increases smoothly from 0 towards  $D = 2$  as the data improves.

At medium values of  $\sigma$ , the prior  $p_*(\vec{y})$  almost ignores the width of the parameter manifold, and cares mostly about its length ( $L_+ = \sqrt{2}/\sigma$  along the diagonal). The behavior of our effective model along this direction is very different to that from Jeffreys prior: in Figure 4B we demonstrate this by plotting the distributions of data implied by these two priors. Jeffreys puts almost no weight near the ends of the long (i.e. stiff, or relevant) parameter's range, because the (sloppy, or irrelevant) width happens to be even narrower there than in the middle. By contrast our effective model puts significant weight on each end, much like the one-parameter model in Figure 4A.

The difference between one and two parameters being relevant (in Figure 3B) is very roughly  $\sigma = 1/7$  to  $\sigma = 1/50$ , a factor 7 in Fisher length, thus a factor 50 in the number of repetitions  $m$  — perhaps the difference between a week's data and a year's. But more realistic models often have manifold lengths spread over many

orders of magnitude [5, 8], and thus have some parameters inaccessible even with centuries of data.

The one-dimensional model along the lower edge of Figure 3A is the effective theory with equal decay constants. This remains true if we allow more parameters  $k_3, k_4, \dots$  in (5), and  $p_*(\vec{y})$  will still place a similar weight there.<sup>8</sup> Measuring  $x_t$  also at later times  $t_3, t_4, \dots$  will add more thin directions to the manifold [7], but the one-dimensional boundary corresponding to equal decay constants will still have significant weight. The fact that such edges give human-readable simpler models (unlike arbitrary submanifolds) was the original motivation for preferring them in [21], and it is very interesting that our optimization procedure has the same preference.

## Discussion

While the three examples we have studied here are very simple, they demonstrate a principled way of selecting optimal effective theories, especially in high-dimensional settings, which we call rational ignorance.

The prior  $p_*(\theta)$  which encodes this selection is the maximally uninformative prior, in the sense of leaving maximum headroom for learning (by Bayesian updating) from data. But its construction depends on  $p(x|\theta)$ , and thus it contains knowledge about the experiment through which we are probing nature. Adapting our effective theory according to the length scales visible to our experiment is an important idea in physics [48], most precisely expressed in the renormalization group [1, 2].

Jeffreys prior  $p_J(\theta)$  also depends on the experiment, but more weakly: it is independent of the number of repetitions  $m$ , precisely because it is the limit  $m \rightarrow \infty$  of the optimal prior [28, 29]. Many experiments would have much higher resolution if they could run for all eternity. The fact that they cannot, that data is finite, is an important restriction on the accuracy of our knowledge.

This limit is where we differ from earlier work in statistics. Bernardo's reference prior [26] maximizes the same mutual information, but always in the  $m \rightarrow \infty$  limit where it gives a smooth analytically tractable function. The use of this mutual information  $I(X; \Theta)$  to quantify what can be learned from an experiment goes back to Lindley [23]. That finite information implies a discrete distribution was known in the engineering literature at least since Färber and Smith [33, 34]. But what has been overlooked is that this discreteness is useful for avoiding the pathologies of Jeffreys prior on the sort of high-dimensional parameter spaces generic in science [5].

<sup>8</sup> If we have more parameters than measurements then the model must be singular. In fact the exponential model of Figure 3 is already slightly singular, since  $k_1 \leftrightarrow k_2$  does not change the data; we could cure this by restricting to  $k_2 \geq k_1$ , or by working with  $\vec{y}$ , to obtain a regular model.

The problem posed by such sloppy manifolds is different to that studied by the information criteria. There, given a large amount of noisy data, one is tempted to fit a model which is more complicated than reality. Avoiding such overfitting improves predictions. The AIC, BIC [14, 49] and related criteria [18, 19, 50–52] for this are derived as subleading terms in the  $m \rightarrow \infty$  limit, in which all (nonsingular) parameters of the true model can be accurately measured. It is precisely by not taking this limit that we retain the distinction between relevant and irrelevant parameters.

Being discrete, the prior  $p_*(\theta)$  is very likely to exclude the true value of the parameter, if such a  $\theta_{\text{true}} \in \Theta$  exists. This is not a flaw: the spirit of effective theory is to focus on what is relevant for describing the data, deliberately ignoring microscopic effects which we know to exist. For example, the same effective theory can emerge from different microscopic physics (as in the universality of critical points describing phase transitions [53]). The relevant degrees of freedom are often quasiparticles (such as the Cooper pairs of superconductivity [54]) which do not exist in the microscopic theory, but give a natural and simple description at the scale being observed. We argued here for this simplicity not on the grounds of the difficulty of simulating  $10^{23}$  electrons, nor of human limitations, but based on the natural measure of information learned.

There is similar simplicity to be found outside of physics, for example in the Michaelis–Menten law for enzyme kinetics [55, 56] which is derived as a limit in which only the ratios of some reaction rates matter, and is useful regardless of the underlying system. In more complicated systems which we cannot solve by hand, and for which the symmetries and scaling arguments used in physics cannot be applied, we hope that our information approach may be useful for identifying the appropriate theory.

### Acknowledgments

We thank Vijay Balasubramanian, William Bialek, Robert de Mello Koch, Jon Machta, James Sethna, Paul Wiggins, and Ned Wingreen for discussion and comments. We thank ICTS Bangalore for hospitality.

H.H.M. was supported by NIH grant R01GM107103. M.K.T. was support by NSF-EPCN 1710727. B.B.M. was supported by a Lewis-Sigler Fellowship and by NSF PHY 0957573. M.C.A. was supported by NCN grant 2012/06/A/ST2/00396.

### Appendix: Algorithms

The standard algorithm for maximizing channel capacity (of discrete memoryless channels) was written down independently by Blahut [46] and Arimoto [45]. This aspect of rate-distortion theory is mathematically the same as the problem we consider, of maximizing mutual information by choosing the prior. The algorithm starts with




Figure 5. **Convergence of the BA algorithm (7).** This is for the one-parameter Gaussian model (2) with  $L = 10$  (comparable to  $m = 10$  in Figure 1). On the right  $\theta$  is discretized into ten times as many points, but  $p_\tau(\theta)$  clearly converges to the same five delta functions.

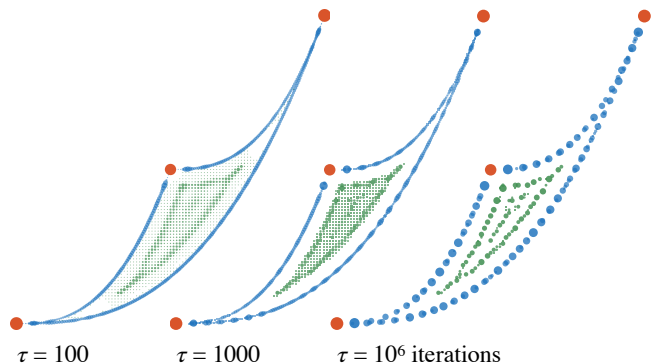


Figure 6. **Convergence of the BA algorithm for the exponential model.** This shows  $p_*(\tilde{y})$  for the case  $\sigma = 1/50$ , with  $\tilde{y}$  discretized on a grid of spacing  $1/100$  in the bulk and  $1/200$  along the boundaries of the allowed region.

$p_0(\theta) = \text{const.}$ , and then at each time step updates this by

$$p_{\tau+1}(\theta) = \frac{1}{Z_\tau} e^{f_{\text{KL}}(\theta)} p_\tau(\theta) \quad (7)$$

where  $Z_\tau = \int d\theta e^{f_{\text{KL}}(\theta)} p_\tau(\theta)$  maintains normalization, and  $f_{\text{KL}}(\theta) = D_{\text{KL}}(p(x|\theta) \| p(x))$  is computed with  $p_\tau(\theta)$ . Since this is a convex optimization problem, the algorithm is guaranteed to converge to the global maximum. This makes it a good tool to see discreteness emerging.

Figures 5 and 6 show the progress of this algorithm for the one- and two-dimensional models in the text. We stress that the number and positions of the peaks which form are unchanged when the discretization of  $\theta$  is made much finer. Notice also that the convergence to delta functions happens much sooner near to the boundaries than in the interior. The convergence to the correct mutual information  $I(X; \Theta)$ , and towards the optimum distribution on data space  $p(x)$ , happens much faster than the convergence to the correct number of delta functions.

Because  $\theta$  must be discretized for this procedure, it is poorly suited to high-dimensional parameter spaces.

However once we know that  $p_*(\theta)$  is discrete, it is natural to consider algorithms exploiting this. With  $K$  atoms, we can adjust their positions  $\vec{\theta}_a$  and weights  $\lambda_a$  using gradients

$$\begin{aligned}\frac{\partial \text{MI}}{\partial \theta_a^\mu} &= \lambda_a \int dx \frac{\partial p(x|\vec{\theta})}{\partial \theta^\mu} \log \frac{p(x|\vec{\theta})}{p(x)} \Big|_{\vec{\theta}=\vec{\theta}_a} \\ \frac{\partial \text{MI}}{\partial \lambda_a} &= f_{\text{KL}}(\vec{\theta}_a) - 1.\end{aligned}\quad (8)$$

Figures 1A, 2A, 4 and the square plot points in Figure 3 were generated this way. This optimization is not a

convex problem (there is some tendency to place two atoms on top of each other, and thus use too few points of support) but it can often find the optimum solution. We can confirm this by calculating  $f_{\text{KL}}(\theta)$  everywhere — any points for which this is larger than its value at the atoms indicates that we do not have the optimal solution, and should add an atom.

Monte Carlo algorithms for this problem have been investigated in the literature, see [57, 58] and especially [59]. (Incidentally, we observe that [57]’s table 1 contains a version of scaling law (4), with  $\zeta \approx 1/2$ . No attempt was made there to use the optimal number of atoms, only to calculate the channel capacity to sufficient accuracy.)

- 
- [1] L. P. Kadanoff, *Scaling laws for Ising models near  $T_c$* , Physics **2** (1966) 263–272.
  - [2] K. G. Wilson, *Renormalization group and critical phenomena. 1. Renormalization group and the Kadanoff scaling picture*, Phys. Rev. **B4** (1971) 3174–3183.
  - [3] J. L. Cardy, *Scaling and renormalization in statistical physics*. Cambridge Univ. Press, 1996.
  - [4] J. J. Waterfall, F. P. Casey, R. N. Gutenkunst, K. S. Brown, C. R. Myers, P. W. Brouwer and J. P. Sethna, *Sloppy-model universality class and the Vandermonde matrix*, Phys. Rev. Lett. **97** (2006) 150601–4.
  - [5] R. N. Gutenkunst, J. J. Waterfall, F. P. Casey, K. S. Brown, C. R. Myers and J. P. Sethna, *Universally Sloppy Parameter Sensitivities in Systems Biology Models*, PLoS Comp. Biol. **3** (2007) e189–8.
  - [6] M. K. Transtrum, B. B. Machta and J. P. Sethna, *Why are Nonlinear Fits to Data so Challenging?*, Phys. Rev. Lett. **104** (2010) 060201 [arXiv:0909.3884].
  - [7] M. K. Transtrum, B. B. Machta and J. P. Sethna, *Geometry of nonlinear least squares with applications to sloppy models and optimization*, Phys. Rev. **E83** (2011) 036701 [arXiv:1010.1449].
  - [8] B. B. Machta, R. Chachra, M. K. Transtrum and J. P. Sethna, *Parameter Space Compression Underlies Emergent Theories and Predictive Models*, Science **342** (2013) 604–607 [arXiv:1303.6738].
  - [9] M. K. Transtrum, B. B. Machta, K. S. Brown, B. C. Daniels, C. R. Myers and J. P. Sethna, *Perspective: Sloppiness and emergent theories in physics, biology, and beyond*, J. Chem. Phys. **143** (2015) 010901 [arXiv:1501.07668].
  - [10] T. O’Leary, A. C. Sutton and E. Marder, *Computational models in the age of large datasets*, Cur. Op. Neurobio. **32** (2015) 87–94.
  - [11] T. Niksic and D. Vretenar, *Sloppy nuclear energy density functionals: effective model reduction*, Phys. Rev. **C94** (2016) 024333 [arXiv:1606.08617].
  - [12] E. Dufresne, H. A. Harrington and D. V. Raman, *The geometry of sloppiness*, arXiv:1608.05679.
  - [13] G. Bohner and G. Venkataraman, *Identifiability, reducibility, and adaptability in allosteric macromolecules*, J. Gen. Physiol. **149** (2017) 547–560.
  - [14] H. Akaike, *A new look at the statistical model identification*, IEEE Trans. Automat. Contr. **19** (1974) 716–723.
  - [15] N. Sugiura, *Further analysts of the data by Akaike’s information criterion and the finite corrections*, Comm. Stat. Th. Meth. **7** (1978) 13–26.
  - [16] V. Balasubramanian, *Statistical inference, Occam’s razor, and statistical mechanics on the space of probability distributions*, Neural Comp. **9** (1997) 349–368.
  - [17] I. J. Myung, V. Balasubramanian and M. A. Pitt, *Counting probability distributions: differential geometry and model selection*, PNAS **97** (2000) 11170–11175.
  - [18] D. J. Spiegelhalter, N. G. Best, B. P. Carlin and A. Van Der Linde, *Bayesian measures of model complexity and fit*, J. Roy. Stat. Soc. B **64** (2002) 583–639.
  - [19] S. Watanabe, *Asymptotic equivalence of Bayes cross validation and widely applicable information criterion in singular learning theory*, JMLR **11** (2010) 3571–3594 [arXiv:1004.2316].
  - [20] P. A. Wiggins and C. H. LaMont, *Information-based inference in sloppy and singular models*, arXiv:1506.05855.
  - [21] M. K. Transtrum and P. Qiu, *Model reduction by manifold boundaries*, Phys. Rev. Lett. **113** (2014) 098701–6.
  - [22] C. E. Shannon, *A mathematical theory of communication*, Bell Sys. Tech. J. **27** (1948) 623–656.
  - [23] D. V. Lindley, *On a Measure of the Information Provided by an Experiment*, Ann. Math. Statist. **27** (1956) 986–1005.
  - [24] A. Rényi, *On some basic problems of statistics from the point of view of information theory*, Proc. 5th Berkeley Symp. on Math. Statist. and Prob. (1967) 531–543. [projecteuclid.org].
  - [25] H. Jeffreys, *An Invariant Form for the Prior Probability in Estimation Problems*, Proc. Roy. Soc. A **186** (1946) 453–461.
  - [26] J. M. Bernardo, *Reference Posterior Distributions for Bayesian Inference*, J. Roy. Stat. Soc. B **41** (1979) 113–147. [www.uv.es/~bernardo/].
  - [27] J. O. Berger, J. M. Bernardo and M. Mendoza, *On Priors that Maximize Expected Information*, Recent Developments in Statistics and Their Applications (1988) 1–20. [www.uv.es/~bernardo/].
  - [28] B. S. Clarke and A. R. Barron, *Jeffreys’ prior is asymptotically least favorable under entropy risk*, J. Stat. Plan. Infer. **41** (1994) 37–60.
  - [29] H. R. Scholl, *Shannon optimal priors on independent*

- identically distributed statistical experiments converge weakly to Jeffreys' prior, *Test* **7** (1998) 75–94.
- [30] J. E. Kerrich, *An Experimental Introduction to the Theory of Probability*. Copenhagen: E Munksgaard, 1946.
- [31] C. O'Lunaigh, *CERN data centre passes 100 petabytes*, [home.cern](http://home.cern) (2013).
- [32] D. Simpson, H. Rue, A. Riebler and T. G. Martins, *Penalising Model Component Complexity: A Principled, Practical Approach to Constructing Priors*, *Statist. Sci.* **32** (2017) 1–28 [[arXiv:1403.4630](https://arxiv.org/abs/1403.4630)].
- [33] G. Färber, *Die Kanalkapazität allgemeiner Übertragungskanäle bei begrenztem Signalwertbereich beliebigen Signalübertragungszeiten sowie beliebiger Störung*, *Arch. Elektr. Übertr.* **21** (1967) 565–574.
- [34] J. G. Smith, *The information capacity of amplitude-and variance-constrained scalar gaussian channels*, *Information and Control* **18** (1971) 203–219.
- [35] S. L. Fix, *Rate distortion functions for squared error distortion measures*, *Proc. 16th Annu. Allerton Conf. Commun. Control Comput.* (1978) 704–711.
- [36] K. Rose, *A mapping approach to rate-distortion computation and analysis*, *IEEE Trans. Inform. Theory* **40** (1994) 1939–1952.
- [37] Z. Zhang, *Discrete Noninformative Priors*. PhD thesis, Yale University, 1994. [UMI 9523257].
- [38] M.-H. Chen, D. K. Dey, P. Müller, D. Sun and K. Ye, *Frontiers of statistical decision making and Bayesian analysis: in honor of James O. Berger*. Springer, New York, NY, 2010.
- [39] C. A. Sims, *Rational Inattention: Beyond the Linear-Quadratic Case*, *American Economic Review* **96** (2006) 158–163.
- [40] J. Jung, J. Kim, F. Matějka and C. A. Sims, *Discrete Actions in Information-Constrained Decision Problems*, . [[princeton.edu/~sims/#RIDiscrete](http://princeton.edu/~sims/#RIDiscrete)].
- [41] S. Laughlin, *A simple coding procedure enhances a neuron's information capacity*, *Z. Naturforsch.* **36 c** (1981) 910–912.
- [42] G. Tkačik, C. G. Callan and W. Bialek, *Information flow and optimization in transcriptional regulation*, *PNAS* **105** (2008) 12265–12270 [[arXiv:0705.0313](https://arxiv.org/abs/0705.0313)].
- [43] M. D. Petkova, G. Tkačik, W. Bialek, E. F. Wieschaus and T. Gregor, *Optimal decoding of information from a genetic network*, [arXiv:1612.08084](https://arxiv.org/abs/1612.08084).
- [44] A. Mayer, V. Balasubramanian, T. Mora and A. M. Walczak, *How a well-adapted immune system is organized*, *PNAS* **112** (2015) 5950–5955 [[arXiv:1407.6888](https://arxiv.org/abs/1407.6888)].
- [45] S. Arimoto, *An algorithm for computing the capacity of arbitrary discrete memoryless channels*, *IEEE Trans. Inform. Theory* **18** (1972) 14–20.
- [46] R. Blahut, *Computation of channel capacity and rate-distortion functions*, *IEEE Trans. Inform. Theory* **18** (1972) 460–473.
- [47] S. H. Strogatz, *Nonlinear Dynamics And Chaos*. Sarat Book House, 2007.
- [48] P. W. Anderson, *More Is Different*, *Science* **177** (1972) 393–396.
- [49] G. Schwarz, *Estimating the dimension of a model*, *Ann. Statist.* **6** (1978) 461–464.
- [50] J. Rissanen, *Modeling by Shortest Data Description*, *Automatica* **14** (1978) 465–471.
- [51] C. S. Wallace and D. M. Boulton, *An Information Measure for Classification*, *The Computer Journal* **11** (1968) 185–194.
- [52] S. Watanabe, *A widely applicable Bayesian information criterion*, *JMLR* **14** (2013) 867–897 [[arXiv:1208.6338](https://arxiv.org/abs/1208.6338)].
- [53] R. W. Batterman, *Philosophical Implications of Kadanoff's Work on the Renormalization Group*, *J. Stat. Phys.* **167** (2017) 559–574. [[pitt.edu/~rbatterman/](http://pitt.edu/~rbatterman/)].
- [54] J. Bardeen, L. N. Cooper and J. R. Schrieffer, *Theory of superconductivity*, *Phys. Rev.* **108** (1957) 1175–1204.
- [55] L. Michaelis and M. L. Menten, *Die Kinetik der Invertinwirkung*, *Biochem Z* **49** (1913) 333–369.
- [56] L. Michaelis and M. L. Menten, *The kinetics of invertin action*, *FEBS Letters* **587** (2013) 2712–2720. [Translation by T. R. C. Boyde].
- [57] C.-I. Chang and L. D. Davisson, *On calculating the capacity of an infinite-input finite (infinite)-output channel*, *IEEE Trans. Inform. Theory* **34** (1988) 1004–1010.
- [58] J. Lafferty and L. Wasserman, *Iterative Markov chain Monte Carlo computation of reference priors and minimax risk*, *Proc. 17th conf. Uncert. AI* (2001) 293–300 [[arXiv:1301.2286](https://arxiv.org/abs/1301.2286)].
- [59] J. Dauwels, *Numerical computation of the capacity of continuous memoryless channels*, *Proc. 26th Symp. Inf. Th. Benelux* (2005). [[www.dauwels.com](http://www.dauwels.com)].