




The relative merit of empirical priors in non-identifiable and sloppy models: Applications to models of learning and decision-making

Empirical priors

Mikhail S. Spektor^{1,2}  · David Kellen³

Published online: 27 March 2018
© Psychonomic Society, Inc. 2018

Abstract

Formal modeling approaches to cognition provide a principled characterization of observed responses in terms of a set of postulated processes, specifically in terms of parameters that modulate the latter. These model-based characterizations are useful to the extent that there is a clear, one-to-one mapping between parameters and model expectations (identifiability) and that parameters can be recovered from reasonably sized data using a typical experimental design (recoverability). These properties are sometimes not met for certain combinations of model classes and data. One suggestion to improve parameter identifiability and recoverability involves the use of “empirical priors”, which constrain parameters according to a previously observed distribution of values. We assessed the efficacy of this proposal using a combination of real and artificial data. Our results showed that a point-estimate variant of the empirical-prior method could not improve parameter recovery systematically. We identified the source of poor parameter recovery in the low information content of the data. As a follow-up step, we developed a fully Bayesian variant of the empirical-prior method and assessed its performance. We find that even such a method that takes the covariance structure of the parameter distributions into account cannot reliably improve parameter recovery. We conclude that researchers should invest additional efforts in improving the informativeness of their experimental designs, as many of the problems associated to impoverished designs cannot be alleviated by modern statistical methods alone.

Keywords Identifiability · Empirical priors · Reinforcement learning · Prospect theory

A key aspect in the development and testing of psychological theories of cognition is the increasing reliance on formal modeling approaches (for an introduction, see Lewandowsky & Farrell, 2010). Through formal models, researchers can characterize the observed data in terms of latent cognitive processes such as attention, memory retrieval, or response

biases. Associated to each of these different processes are parameters that determine their expression and role: For example, consider Schweickert’s (1993) model of short-term memory retrieval illustrated in Fig. 1. This model postulates a parameter I quantifying the probability that the representation of a studied word in short-term memory is intact and a parameter R quantifying the probability that in the absence of an intact representation the word can be successfully reintegrated. According to this model, the probability of an item being correctly recalled is $I + (1 - I) \times R$ (item is intact or can be successfully reintegrated), whereas the probability of an incorrect recall is $(1 - I) \times (1 - R)$ (item is not intact and reintegration is not successful).

When specifying the multiple components within a given model, researchers face several challenges. Among them is the need to ensure that model parameters are *identifiable* (Bamber & van Santen 1985; Moran, 2016). Broadly speaking, identifiability concerns the notion that each combination of parameter values (e.g., I and R) in a model

✉ Mikhail S. Spektor
mikhail@spektor.ch

David Kellen
davekellen@gmail.com

¹ Faculty of Psychology, University of Basel, Basel, Switzerland

² Department of Psychology, University of Freiburg, Freiburg, Germany

³ College of Arts and Sciences, Syracuse University, Syracuse NY, USA




Fig. 1 Schweickert's (1993) redintegration model

yields a unique set of expectations. When the parameters of a model are identifiable, the modeler can be sure that there is a unique set of parameters providing the best match between the model's expectations and the data. The only limitation then is the informativeness of the data. Unfortunately, this is not the case in Schweickert's model, as its parameters are not identifiable: The same recall probabilities can be obtained by trading off I and R . For example, a correct-recall probability of .90 can be obtained with an infinite range of $\{I, R\}$ pairs. From $\{I = .90, R = .00\}$, which states that items are very likely to have an intact representation but no hope for redintegration, to $\{I = .00, R = .90\}$, which states that there are no intact memory representations but redintegration is nevertheless highly likely. In order to address the non-identifiability of these parameters, researchers have relied on complex experimental designs through which principled constraints can be imposed (Hulme, Roodenrys, Schweickert, Brown, et al. 1997; Buchner & Erdfelder, 2005; Schweickert, 1993). In general, the identifiability of parameters will be a function of several factors, such as the model class, number of parameters to be estimated, parametric assumptions made, the data to which the models are fit, and the method used to fit the models (Bamber & van Santen, 1985; Batchelder & Riefer, 1990; Moran, 2016; Ahn, Krawitz, Kim, Busemeyer, & Brown, 2011; Wetzels, Vandekerckhove, Tuerlinckx, & Wagenmakers, 2010).

The importance of parameter identifiability can hardly be overstated, given that it ensures that the theoretically motivated characterization of the data yielded by a model is unique. In lay terms, we want to make sure that the model tells us a unique story for a given set of data, not a multitude of distinct stories. After all, the parameters of cognitive models are supposed to reflect psychologically meaningful processes. In light of the importance of parameter identifiability, many approaches have been developed for overcoming difficulties with it. The present work is

concerned with one specific approach for overcoming such issues: The use of parametric approximations to the distributions of parameter values obtained from a separate dataset as informative priors, and incorporating them in subsequent modeling efforts.

In what follows, we will first provide an overview of the different kinds of identifiability and the notion of *sloppiness*. We will then discuss general approaches of alleviating the problems associated with sloppiness and non-identifiability and focus on one specific method, the *empirical-prior approach* (Gershman, 2016). We evaluate the point-estimate variant of the empirical-prior approach as proposed by Gershman (2016) in a well-known class of reinforcement-learning models. We assess the improvements brought about by this approach in comparison to the improvements coming from simple extensions of the experimental design. Finally, we develop a Bayesian variant of the empirical-prior approach and apply it to another well-known class of models, this time concerned with *decision-making under risk*.

Foreshadowing our results, we found that the empirical-prior approach cannot recover the true parameter population distributions, does not improve parameter recovery, and is fragile to mismatches between the true parameter population distributions and the prior used. Similar results were obtained with a Bayesian variant of the empirical-prior approach. In contrast with these rather disappointing results, small changes in the experimental design showed clear improvements in parameter recovery. We conclude that for researchers interested in inferring the ground truth from empirical observations, the shortcomings of impoverished experimental designs that fail to constrain parameter estimates cannot be compensated for through the use of statistical methods such as the empirical-prior approach. We highlight the differences between this so-called “question of inversion” and the “question of inference” that is concerned with coherent interpretation of the data irrespective of the ground truth that generated them.

Overcoming varieties of non-identifiability and sloppiness

A more nuanced view of identifiability is given by the distinction between *global identifiability*, which concerns the identifiability of the parameters irrespective of any particular data within a given experimental design, and *local identifiability*, which is only concerned with identifiability of a model for a particular set of data (for a detailed discussion and examples, see Schmittmann et al., 2010). A model can be globally identifiable but locally non-identifiable due to sparse data (e.g., several empty cells in a multinomial distribution) and/or extreme performance

(e.g., ceiling or floor effects). For example, a model that is constrained to predict either the sequence a-a-a-a or b-b-b-b under different parameter values will be locally non-identifiable when observing the sequence a-b-a-b, as both sequences the model is able to produce account for the data equally well/badly.

A concept closely related to identifiability is “sloppiness” (Brown & Sethna, 2003). Although identifiable, the parameters of a “sloppy” model can be adjusted to *partially* compensate any change in the expectations produced by the variation of another parameter. One consequence of sloppiness is that it becomes difficult to determine the parameter values of the underlying data-generating process. These difficulties can be demonstrated in parameter-recovery simulations, in which artificial data are generated with known parameter values. The very same model is then used to fit the data, so that the original and estimated parameter values can be compared. Note that non-identifiability corresponds to the worst-case scenario in terms of sloppiness, with parameters being able to *perfectly* compensate for the changes in other parameters.

Given the importance of parameter estimates in a model-based characterization of behavioral phenomena, it follows that parameter identifiability is only a minimal and insufficient requirement. After all, a model can have identifiable parameters but nevertheless manifest poor parameter recovery under a more realistic setting. In line with this notion, considerable efforts have been made to identify and ameliorate difficulties with parameter recovery. For example, White, Servant, and Logan (2017) showed that some of the parameters in drift-diffusion models for conflict tasks, although identifiable, have poor recoverability. In order to mitigate these issues, White et al. (2017) suggested the use of derived measures that try to make up for the parameter trade-offs they observed. In other domains (e.g., decision-making under risk), alternative solutions such as parameter restrictions have been proposed: Nilsson, Rieskamp, and Wagenmakers (2011) evaluated the value function $v(\cdot)$ of prospect theory (Kahneman & Tversky, 1979), which assumes that the subjective representation of monetary gains $x > 0$ follows $v(x) = x^{\alpha^+}$, whereas for losses $x < 0$ it follows $v(x) = -\lambda|x|^{\alpha^-}$. Nilsson et al. showed that the loss-aversion parameter λ and the diminishing-sensitivity parameter for losses α^- are extremely hard to recover as they tend to serve very similar purposes and can therefore trade off with little to no cost. Their solution to this parameter-recovery problem was to simplify the model by setting α^- to be equal to its gain-outcome counterpart α^+ . But as discussed later on, other issues in parameter recovery remain to be addressed.

A more general modeling approach that can address some of the difficulties in parameter recoverability consists of relying on *hierarchical* or *random-effect* implementations

of models. A key aspect of these implementations is that the estimations of individual parameters are informed by the overall sample, capitalizing on the similarities across individuals by shrinking individual estimates towards a central tendency of the group level and thus preventing extreme, noise-driven parameter estimates. This approach is particularly helpful if the experimental design is informative and the model identifiable, but there are not enough data per individual to estimate their respective parameters with high precision (Broomell & Bhatia, 2014). In the cognitive-modeling literature, hierarchical models are typically fitted using Bayesian parameter estimation (e.g., Katahira, 2016; Steingroever, Wetzels, & Wagenmakers, 2014; Wetzels et al., 2010; Ahn et al., 2011).

In contrast to the conventional maximum likelihood estimation (MLE) methods often used in model fitting, Bayesian approaches require the specification of *prior distributions* for each of the parameters, representing the (prior) beliefs that parameters will take on certain values. These priors are updated using the information provided by the data, resulting in posterior distributions. MLE does not require such a prior nor does it yield a distribution—only the set of parameter values for which the likelihood of the data is maximal. Whether a model is fitted using MLE or Bayesian parameter estimation affects the assessment of identifiability, sloppiness, and parameter recovery. Non-identifiability and sloppiness will lead to regions of the joint posterior distributions which have equal and near-equal density, respectively. When using MLE, trade-offs between sloppy models’ parameters can be observed in the covariance matrix of parameter estimates (see Li, Lewandowsky, & DeBrunner, 1996). When using a Bayesian parameter-estimation framework, parameter trade-offs are reflected in the covariances of posterior samples and the (multivariate) posterior distributions of parameters. Specifically, they often manifest themselves as ridges in the joint posterior distributions (e.g., Scheibehenne & Pachur, 2015, Fig. 1).

A compromise between MLE and Bayesian estimation is offered by the *maximum a-posteriori* (MAP) method (Cousineau & Hélie, 2013). MAP introduces prior parameter distributions that are used to weight the model’s likelihood function. This prior-weighted MLE yields the modes of the posterior parameter distributions that would be obtained with a fully Bayesian approach using the same set of priors. Both the fully Bayesian approach and MAP have been alluded to as ways to attenuate parameter-identifiability problems (e.g., Moran, 2016) by using an *informative prior*. An informative prior carries information about the parameters that goes beyond the data at hand, such as from previous empirical work or from theoretical considerations (for an overview, see Lee & Vanpaemel, 2017). For instance, if one has good reasons to believe that reasonable

parameter values should most often lie within a certain range, one can specify a prior that overweights that specific range relative to the rest. This weighting would discourage estimates to go outside this expected range unless the data strongly support that. Hierarchical parameter estimation poses a special case of the use of informative priors in which the parameter estimates of individuals inform each other.

Gershman's (2016) empirical-prior approach

Gershman (2016) recently argued that one way to obtain an informative prior is to use the distribution of MLE estimates obtained from a separate dataset in an attempt to approximate the population distribution of parameter values—an *empirical prior*. In the context of reinforcement-learning models (Sutton & Barto 1998), Gershman demonstrated that the MAP method, together with empirical priors, could improve model performance in several ways: more reliable parameter estimates, improved characterization of individual differences, and increases in predictive accuracy.

The improvements reported by Gershman's (2016) empirical-prior approach in the context of reinforcement-learning modeling are quite fortunate, as they directly address some long-standing challenges in this domain. Reinforcement-learning models are regularly adopted as a way to analyze repeated trial-and-error decisions in psychology and neuroscience (e.g., Yechiam & Busemeyer, 2005; Schulze, van Ravenzwaaij, & Newell, 2015; Erev & Barron, 2005; Baron & Erev, 2003; Niv et al., 2015; Dayan & Daw, 2008; Chase, Kumar, Eickhoff, & Dombrovski, 2015; Dayan & Balleine, 2002). Despite their prominence, these models have well-documented cases of parameter non-identifiability and sloppiness (e.g., Humphries, Bruno, Karpievitch, & Wotherspoon, 2015; Wetzels et al., 2010; but see, e.g., Ahn et al., 2011, 2014; Steingroever, Wetzels, & Wagenmakers, 2013, for examples of satisfactory parameter identifiability). One illustrative example was recently given by Humphries et al. (2015), who fitted a popular reinforcement-learning model to choice data obtained with the Iowa gambling task (Bechara, Damasio, Damasio, and Anderson, 1994). Humphries et al. showed that the best fits could often be achieved under very different sets of parameter values that yielded quite distinct accounts of the data. For example, one participant had a set of best-fitting parameters indicating that s/he had good memory and produced impulsive, consistent choices. However, another equally good set parameter estimates for that same participant suggested that s/he could be characterized as an individual with inconsistent choices and poor memory, and who focused more on losses than wins (p. 24). Reducing the model complexity by restricting the number of free parameters limited

the degree to which parameters traded off, but also limited the richness of the characterization provided by the model.

Despite its purported advantages, some aspects of Gershman's (2016) empirical-prior approach require further scrutiny. First, it is not clear *how* this approach can mitigate the problems of parameter identifiability and recoverability. Since it relies on marginal prior parameter distributions that are independent from each other, the constraints imposed do not extend to the way parameters can jointly vary to produce equivalent or very similar results. In somewhat broad strokes, it can be said that Gershman's empirical-prior approach is attempting to tackle a problem of parameter *covariance* by constraining *variances*. It could very well be that the improvements reported by Gershman (2016) result from shifting the parameter estimates towards values that are more commonly observed in the population, thus avoiding overfitting.

Second, Gershman (2016) approach assumes that the empirical priors obtained are somewhat reasonable approximations to the distributions of parameter values in the population. Given that the reinforcement-learning models considered in his application suffer from identifiability and recoverability issues, it seems unlikely that such an approximation is achieved to any reasonable degree. If the empirical priors are themselves based on unreliable parameter estimates that do not match the actual data-generating parameters, it is not clear how they could contribute to mitigating any identifiability or recoverability problems. In short, the approach seems to be stuck with a "chicken-or-the-egg" type of problem.

Evaluating the empirical-prior approach

Given the standing questions regarding Gershman's (2016) empirical-prior approach, we conducted additional evaluations. These evaluations required knowledge of the true parameter values that generated any given dataset, something that is only possible when the data are artificially generated. Therefore, we implemented a series of simulation studies, through which we tested whether that empirical-prior approach improves the recovery of parameter estimates in the reinforcement-learning models considered by Gershman, and how these estimates are affected by the (mis)match between the empirical priors and the actual distribution of parameters in the population. We will later extend these simulation studies to the domain of decision-making under risk.¹

¹All our simulation results and analysis scripts are provided on the Open Science Framework at <https://osf.io/2ws78/>.




Fig. 2 Example of a typical probability-learning task

Method

Data

We used the data from Gershman (2015) as the basis for our simulations. These data came from a probability-learning task in which participants were presented with two options differing in terms of their probabilities of yielding a reward (0 or 1; see Fig. 2 for an example of such a probability-learning task). Participants were provided with partial feedback such that they were only informed of the outcome yielded by the chosen option. Each participant engaged in four blocks of 25 trials each, for a total of 100 trials. Their responses came from four experiments, which only differed in terms of the reward probabilities used, for a total of 166 participants.²

Models

We focused our comparisons on the same four reinforcement-learning models that Gershman (2016) considered, all of which conform to the basic structure of a *Q-learning model* (Sutton & Barto, 1998). According to this model, participants keep subjective expectations Q for each of the options and update them based on the feedback they receive from the choices. The updating mechanism typically used is *temporal-difference learning*, in which the expectation regarding an option X is updated based on the difference between the observed reward $R(X)$ and an expected reward $Q(X)$. This difference is conventionally referred to as the *reward prediction error*. Formally, the

subjective expectation $Q_{t+1}(X)$ for option X on trial $t + 1$ is given by

$$Q_{t+1}(X) = Q_t(X) + \eta (R_t(X) - Q_t(X)), \quad (1)$$

where $R_t(X)$ is the reward coming from option X at trial t , and η is the learning rate. Note that the difference in parentheses is the reward prediction error on trial t . Initial expectations $Q_0(\cdot)$ are set to 0 and on every trial the probabilities of the choices are governed by a logistic function with scaling parameter β .³ The simplest model in our set, \mathcal{M}_1 , assumes a single learning rate η . Model \mathcal{M}_2 extends \mathcal{M}_1 by assuming different learning rates η for the different reward prediction errors:

$$\eta = \begin{cases} \eta^+ & \text{if } R_t(X) - Q_t(X) \geq 0 \\ \eta^- & \text{if } R_t(X) - Q_t(X) < 0 \end{cases}. \quad (2)$$

The third model, \mathcal{M}_3 , also builds on \mathcal{M}_1 by including a *stickiness* mechanism that attempts to capture the notion that individuals tend to repeat previous choices independently of the outcome obtained in the previous trial (also referred to as “perseveration” or “choice inertia”; Yechiam & Ert, 2007; Worthy, Pang, & Byrne, 2013). Formally, c_{t-1} , the choice made on trial $t - 1$, introduces an additive constant ω to the expectation of this option on trial t :

$$Q_t(X) = Q_{t-1}(X) + \eta (R_{t-1}(X) - Q_{t-1}(X)) + \omega, \text{ if } c_{t-1} = X. \quad (3)$$

Note that $\omega < 0$ penalizes repeated choices of the same option, and $\omega > 0$ encourages repeated choices from the same option, irrespective of the rewards it yields. The final model, \mathcal{M}_4 , combines both the two learning rates introduced in \mathcal{M}_2 with the stickiness parameter included in \mathcal{M}_3 .

Altogether, the four models have the following parameters:

- \mathcal{M}_1 : η and β
- \mathcal{M}_2 : η^+ , η^- , and β
- \mathcal{M}_3 : η , β , and ω
- \mathcal{M}_4 : η^+ , η^- , β , and ω

Priors

We obtained empirical priors for each of the models by fitting them to the individual datasets using a differential-evolution algorithm as implemented in `DEoptim` (Mullen, Ardia, Gil, Windover, & Cline, 2011) for R (R Development Core Team, 2008). Following Gershman’s (2016) procedure, we restricted the unbounded parameters β and ω to very broad but not impossible ranges ($[0, 50]$ and $[-5, 5]$,

³We explored the influence of using different $Q_0(\cdot)$ values on parameter recovery. We found slightly improved parameter recovery in case of $Q_0(\cdot) = 0.5$, but only if the ground-truth $Q_0(\cdot)$ was also 0.5. In case of a mismatch in either direction (i.e., data were generated with $Q_0(\cdot) = 0.5$ but fit with $Q_0(\cdot) = 0$ or vice versa), parameter recovery dropped dramatically.

²One of the participants completed only three instead of four blocks.

respectively). We used the algorithm with mostly default settings, but increased the number of population members to 50 and the maximum allowed population generations to 100. The population members were initialized randomly within the parameter boundaries.

To facilitate sampling from the empirical prior distribution without committing to too many auxiliary assumptions, we fitted the observed parameter estimates with *Gaussian mixture models* (GMMs).⁴ These mixtures were obtained by first linearly transforming all parameters to the unit scale [0,1] and then applying an inverse-probit transformation so that they would be represented along the real line.⁵ We then fitted mixture models with up to ten component or base distributions and selected the best-performing mixture model using the Bayesian information criterion (BIC; Schwarz, 1978). The BICs of the fitted GMMs are reported in Table 4 in the Appendix.

In addition to the empirical priors, we also considered uniform prior distributions that simply reflect the parameter bounds used in the fitting procedure: [0, 1], [0, 50], and [−5, 5] for the learning rates, scaling parameter, and stickiness, respectively. The results obtained with these uniform priors provide us with the yardstick with which we can evaluate the benefits associated with the use of empirical priors.⁶

Simulation procedure

We began by generating 1,000 sets of 4×25 reward sequences based on the reward probabilities associated to the different choice options. Afterwards, we drew 1,000 independent parameter-set samples for each model \times prior combination. These draws were used to create the so-called *empirical populations* (as they were obtained from the empirical priors) and *uniform populations* (obtained from the uniform priors). The sampled parameters were then used to simulate responses for all of the generated reward sequences.

Parameter recovery

To assess parameter recovery, we fitted the simulated individual responses twice, once using MLE and once again using MAP in conjunction with the empirical priors. As

⁴The use of GMMs deviates from Gershman's (2016) procedure. The reason for this deviation is that the parametric distribution families he adopted were not able to capture the multimodalities found in the parameter estimates we obtained.

⁵To avoid non-finite values on the real scale, we truncated values on the unit scale at 10^{-10} and $1 - 10^{-10}$.

⁶Note that the use of MAP in conjunction with uniform priors for all parameters yields identical results to conventional MLE when constrained by the same boundaries.

before, we relied on a differential-evolution algorithm. To facilitate model fitting, we initialized the population members of the algorithm with samples from the ground-truth population distributions. Parameter recovery was assessed by regressing the estimated parameter values with the true data-generating parameters. This was done separately for each of the parameters, models, and population distributions. Within this context, we chose to use the explained variance statistic r^2 as a measure of parameter recovery. The reason is that it quantifies the ability of one parameter to capture the variability found in the estimates obtained from the data.

Results

Empirical priors

As previously discussed, we obtained our empirical priors by fitting GMMs to the distributions of parameter estimates. The best-fitting estimates are reported in Table 5 on a probability scale. To obtain values on the real scale, transformations have to be applied: specifically, for the parameters on the unit [0, 1] scale (learning rates η , η^+ , and η^-), values have to be probit-transformed. For parameters on the [0, 50] scale (sensitivity β), values have to be probit-transformed and then multiplied by 50. For the stickiness parameter ω , values have to be probit-transformed, multiplied by 10, and then have 5 subtracted.

- \mathcal{M}_1 : For the η parameter, the best-fitting solution consists of a mixture of four component distributions. The resulting empirical prior is characterized by a strong bimodality at the edges of the parameter space, with relatively little density in between (see Fig. 3, top row, first column). For the β parameter, the best-fitting solution is also a mixture of four Gaussian distributions. Most of the density is concentrated at the region between 0 and 10 on the real scale, with another peak at the upper boundary of the parameter space (50). Almost no density is found between these peaks (see Fig. 3, top row, third column).
- \mathcal{M}_2 : For the η^+ and η^- parameters, the best-fitting solution is a mixture of three Gaussian distributions. In both cases, the priors are characterized by strong bimodalities at the edges of the parameter spaces, with comparatively little density in between (see Fig. 3, second row, first and second columns). For the β parameter, the best-fitting solution is a mixture of six Gaussian distributions. Most of the density is concentrated at the region between 0 and 10 on the real scale, with another peak at the upper boundary of the parameter space—again, with almost no density in between (see Fig. 3, second row, third column).




Fig. 3 Distribution of individual-level maximum likelihood parameter estimates and the fitted Gaussian mixture models for each of the four reinforcement-learning models. *Empty cells* indicate parameters not present in the respective model. Each cell depicts the distributions of the estimated parameters (*shaded area*) and the fitted mixture model (*black line*). See “Models” for model specifications

- \mathcal{M}_3 : For the η and β parameters, the best-fitting solutions correspond to a mixture of three Gaussian distributions. As with models \mathcal{M}_1 and \mathcal{M}_2 , they are characterized by large peaks at the boundaries of their respective parameter spaces and little density in between (see Fig. 3, third row, first and third columns). Parameter ω , on the other hand, has a highly peaked trimodal distribution (captured by a mixture of six Gaussian distributions). The peaks are at the boundaries (-5 and +5) and at 0 (see Fig. 3, third row, last column).
- \mathcal{M}_4 : The parameter distributions are comparable to the other models’ distributions. For the learning rates η^+ and η^- , the best-fitting solutions are mixtures of four Gaussian distributions. The distribution of β estimates is best captured by a mixture of two Gaussian distributions, and the stickiness parameter ω is best described by a mixture of six Gaussian distributions. In all cases, the distributions are multimodal with peaks at the boundaries of the parameter spaces (and at 0 for ω) and very little density in between (see Fig. 3, bottom row).

Simulation results

Recovery of population parameter distributions We assessed whether the parameter estimates obtained from the simulated data resembled the true population distribution. To do so, we first created 100 equally spaced bins that covered the entire parameter range for each of the parameters. Afterwards, we computed the proportion of the ground-truth parameter values falling within each bin (expected frequencies). We then computed a discrepancy statistic using the sum of squares between the expected frequencies and the binned frequencies of the fitted parameter estimates (observed frequencies). For 10,000 samples of 1,000 parameters from each of the ground-truth parameter distributions, we calculated their discrepancy statistics with respect to the expected frequencies, thus obtaining a distribution of discrepancies. Using the relative rank (RR) of observed frequencies within the distribution of discrepancies, we can calculate the probability $P_R = .50 - |(RR - .50)|$ of such a rank being observed when assuming that the recovered parameters stem from the true population distributions.

Recovery of the parameter distributions was found to be somewhat poor across all models, population distributions, and estimation procedures. For \mathcal{M}_1 , only one out of eight parameter distribution could be recovered, namely the distribution of β stemming from the empirical-prior population when using MLE ($P_R = .16$, all other $P_{RS} \leq .04$; see Fig. 4 for an example illustrating the poor distribution recovery). For \mathcal{M}_2 , the pattern looks worse than for \mathcal{M}_1 , as none of the distributions could be successfully recovered (all $P_{RS} < .01$). For \mathcal{M}_3 , the pattern changes slightly: When using uniform priors, an empirical parameter population distribution can successfully be recovered for the β and ω parameters (both $P_{RS} \geq .12$). Otherwise, no parameter distributions were recovered (all $P_{RS} \leq .05$). For \mathcal{M}_4 , no parameter distributions were recovered (all $P_{RS} \leq .02$).

Individual parameter recovery Parameter recovery of individual parameters, quantified here in terms of r^2 , were generally poor (see the “Overall” column in Table 1). The only exception was the simplest model, \mathcal{M}_1 , which had the best general parameter recovery. \mathcal{M}_2 had the worst parameter recovery for all parameters. The most complex model, \mathcal{M}_4 , is the second-worst in terms of recoverability, followed by \mathcal{M}_3 .

Sensitivity to misspecification To assess the sensitivity to prior misspecification in MAP, we compared the parameter recoveries when the population distributions matched the

empirical priors with the analogous recoveries when the two distributions did not match (e.g., fitting data generated from a uniform parameter population when using the empirical priors). The r^2 values for all models, estimation methods, and priors are reported in Table 1.

For model \mathcal{M}_1 , matching the priors to the true underlying population distribution played an important role. The differences in r^2 between the matching and mismatching priors could be as high as .10. For model \mathcal{M}_2 , failing to match the prior used in MAP to the underlying data-generating parameter distributions led to mixed results. In four cases, recoverability became worse when the two distributions mismatched, but it actually *improved* in two other cases. Turning to model \mathcal{M}_3 , we found a pattern similar to the other models: Matching the parameter distributions was important, yet a mismatch also led to substantially improved parameter recovery in one case. Finally, for model \mathcal{M}_4 , except for the learning rates stemming from the empirical-prior distribution, MLE was in all cases better in recovering the true individual-level parameters.

Interim summary

We explored the merit of using MAP in conjunction with empirical priors as a way to improve parameter recovery in reinforcement-learning models. Using simulations, we found that no method yielded satisfactory results for




Fig. 4 Parameter and distribution recovery for model \mathcal{M}_1 (the best-faring case for both types of recoveries). The axes show the true parameters of the model and the corresponding recovered parameters. The true parameters stem from the empirical-prior distribution, and recovered parameters have been obtained using maximum a-posteriori estimation in conjunction with empirical priors

Table 1 Parameter recoverability (in r^2)

Model	Parameter	Uniform population		“Empirical” population		Overall
		MLE	MAP	MLE	MAP	
\mathcal{M}_1	η	.71	.68	.30	.40	.45
	β	.29	.28	.28	.38	.39
\mathcal{M}_2	η^+	.23	.11	.17	.29	.18
	η^-	.72	.66	.28	.24	.39
	β	.17	.15	.23	.21	.20
\mathcal{M}_3	η	.28	.41	.41	.49	.37
	β	.22	.13	.41	.44	.32
	ω	.72	.64	.43	.42	.48
\mathcal{M}_4	η^+	.28	.21	.28	.35	.26
	η^-	.02	.01	.07	.13	.05
	β	.17	.13	.42	.32	.33
	ω	.73	.43	.44	.22	.45

See “Models” for model specifications. *MLE* = maximum likelihood estimation, *MAP* = maximum a-posteriori estimation, *Uniform population* = uniform group-level distribution, “*Empirical*” *population* = group-level distribution stemming from the prior that is used for MAP

any of the criteria we used (i.e., distribution recovery and individual-parameter recovery), with no method being consistently superior to the other across all models. Parameter recoverability was generally poor and alarmingly so in some cases, raising serious questions on the ability to draw conclusions about underlying psychological processes under the experimental design considered by Gershman (2015). In the hopes of following up this rather negative state of affairs with a more positive message, we explored different ways to improve the present experimental design.

Exploring ways to improve recoverability

In an attempt to improve recoverability, we considered different ways in which the experimental design used by Gershman (2015) could be improved. To keep things as simple as possible, we restricted ourselves to model \mathcal{M}_1 . Also, instead of using either MLE or MAP, we adopted a fully Bayesian approach in which posterior distributions of parameters are obtained. In contrast to the point estimates yielded by MLE and MAP, these posterior distributions can be conveniently used to assess the degree of uncertainty surrounding each parameter estimate. Diffuse posteriors are expected when parameters are not identifiable or sloppy. Note that in some cases, non-identifiability can lead to multimodalities in the marginal posterior distributions, and ridges in the joint posterior distributions (with each ridge reflecting a specific parameter trade-off).

Method and Results

We obtained the posterior distributions using a No-U-turn sampler (Hoffman & Gelman, 2014) as implemented in *Stan* (Carpenter et al., 2017) via the *PyStan* interface (Stan Development Team, 2016a). We ran four randomly initialized chains in parallel for 1,000 total iterations, out of which 500 were used as a warm-up period to tune the sampler’s parameters. These warm-up samples were discarded afterwards. The remaining 500 iterations from each chain were concatenated, resulting in a total of 2,000 samples. We restricted β to be between 0 and 50, just like with the point-estimate fits beforehand. In these analyses, we focused on the range of each parameter’s 95% central posterior interval, divided by the range of its support. The resulting *coverage ratio* yields values between 0 and 1, with 0 indicating that all posterior mass in a single point, and with values approaching 1 indicating that any permissible value is likely (i.e., the data are not informative for the estimation of a given parameter).

Baseline

To ensure a more comprehensive assessment, we engaged in a systematic exploration of the parameter space using a grid search. As our baseline for quantitative comparisons, we used a probability-learning task with 4 blocks of 25 trials each. Within the blocks, the virtual participants chose between two options with reward probabilities (.1, .3), (.2, .4), (.6, .8), and (.7, .9), respectively. We created a linearly spaced 101×101 grid of sensitivities β and learning rates η between 0 and 50 and 0 and 1, respectively. As both η and β values of 0 lead to completely random choices, we excluded these values from the simulations, resulting in a final 100×100 grid. For each of the parameter combinations, we simulated the outcome sequences and responses of ten virtual participants, for a total of 100,000 response vectors. Results for the baseline design and all variants are reported in Table 2.

Compared to the previous simulation results reported in the Individual parameter recovery section, the baseline design showed a generally better parameter recoverability due to the change from MLE to the means of the respective posterior distributions (see Table 2). But as the coverage ratios show, the uncertainty surrounding the estimates was still unsatisfactory: In the case of η , we can only hope to reliably distinguish between extremely low and high learning rates. Parameter β does not even lend itself to such hopes. Note that one critical difference between the two estimation procedures is that in the case of multiple maxima, MLE will only yield one of them, whereas using the mean of the posterior distribution effectively averages across these multiple modes. What these results show is that, if anything,

Table 2 Recoverability of baseline analysis and different manipulations of experimental design (model \mathcal{M}_1)

Parameter	Design	Recoverability		Coverage ratio		
		r^2	Δr^2	M	Md	SD
η	Baseline	.83		.42	.41	.16
	↑Trials	.91	.08	.30	.30	.14
	↑Options	.88	.05	.34	.32	.15
	+Full feedback	.90	.07	.30	.28	.17
β	Baseline	.52		.67	.73	.20
	↑Trials	.62	.10	.58	.66	.24
	↑Options	.61	.09	.60	.67	.23
	+Full feedback	.64	.12	.58	.62	.21

See “Models” for model specification. η =learning rate. β =sensitivity. r^2 = shared variance of true parameters in the grid search and posterior means of the corresponding parameter estimates. Δr^2 = difference between the design change and the baseline (i.e., improvement of parameter recoverability). *Coverage ratio* = range of the 95% Bayesian central posterior interval divided by the range of the support (50 and 1 for β and η , respectively). *Baseline* = 4 blocks of 25 trials each comprising 2 options to choose from. ↑ trials = Variant 1 with 50 instead of 25 trials per block. ↑Options = Variant 2 with 4 options instead of 2 within each block. +Full feedback = Variant 3 where forgone outcomes are presented in addition to obtained outcomes

one is better off by adopting a fully Bayesian approach with non-informative priors than introducing empirical priors over point estimates.

Variant 1: Increase number of trials

We explored how an increase of the number of trials within blocks improves identifiability (Table 2, ↑Trials). We increased the number from 25 (baseline) to 50, leading to a total of 200 instead of 100 trials (baseline). As expected, increasing the number of trials within each block improves recoverability for both parameters, although the coverage ratios still indicate a considerable degree of uncertainty.⁷

Variant 2: Increase number of options

As a second variant, we explored the possibility of increasing the number of options for participants to choose from (while keeping the number of blocks and trials per block constant; Table 2, ↑Options). We formed four blocks of four options with reward probabilities (.1, .2, .3, .4), (.6, .7, .8, .9), (.2, .3, .5, .6), and (.5, .6, .8, .9), respectively. The most notable difference between this variant and the

⁷As a robustness check, we tried increasing the number of trials from 25 per block in the baseline to 500 trials per block. Despite this 20-fold increase in number of trials (resulting in 2,000 trials in total), the coverage ratios, especially for β , did not reach a satisfactory level ($M = .33$, $Md = .37$, $SD = .23$).

first one is that comparable improvements in recoverability are achieved without an increase of the total number of trials.

Variant 3: Provide full feedback

As the last variant, we explored how providing participants with full feedback (i.e., giving feedback about the forgone outcomes) influenced parameter recovery (Table 2, +Full feedback). We assumed that the learning rate is identical for both the chosen and the non-chosen options. Similar to Variant 2, this change of design does not lead to an increased number of observations. Yet, it descriptively provides the greatest improvement of parameter recovery, although the recoverability of β , in particular the coverage ratio, is still somewhat disappointing, as it covers more than half the parameter range on average.

Generalizing the evaluation of the empirical-prior approach: An application to risky-choice modeling

It is possible that our disappointing results with the empirical-prior approach were due to the reliance on point estimates, together with the specific reinforcement-learning models and experimental designs considered by Gershman (2016). In order to evaluate this possibility, we developed a fully Bayesian implementation of the empirical-prior approach, and applied it to a different model class and experimental paradigm.

The basic idea of the fully Bayesian empirical-prior approach is a straightforward extension of the previously used point-estimate empirical-prior method: Instead of fitting GMMs to the point estimates of an MLE-based procedure, the GMMs are fitted to the pooled individual-level posterior distributions. This extension offers two main advantages: First, uncertainty about the parameter estimates used to obtain the empirical priors is directly reflected in the empirical priors obtained, and second, because there are many more data points available per individual, it becomes feasible to estimate the covariance matrices associated with each of the multivariate component distributions in a mixture.

The uncertainty associated with any parameter estimate under a Bayesian framework is directly expressed in that parameter’s posterior distribution. We can use this feature to establish an alternative way of assessing parameter recovery. In addition to computing r^2 and coverage ratios, we now also consider $P(95\%CI)$: The proportion of times that the true parameter was included in the 95% *credible interval* estimated from the generated data. These intervals encompass the central 95% of their respective

parameter's posterior distribution. Ideally, one would expect these credible intervals include the true parameter values with probability .95.

Prospect theory and the risky-choice paradigm

One of the most widely used paradigms in the decision-making literature is the *risky-choice paradigm*. In this paradigm, an individual is requested to express her/his preferences between different options that yield monetary outcomes with known probabilities (decision-making under risk), such as the lottery $A = \begin{pmatrix} \$100 & -\$20 \\ .50 & .50 \end{pmatrix}$ that yields a \$100 gain with probability .50, otherwise a \$20 loss, and an option $B = \begin{pmatrix} \$80 & \$0 \\ .50 & .50 \end{pmatrix}$ that yields a \$80 gain with probability .50, otherwise nothing. Individuals' preferences regarding options of this kind are expected to capture their subjective representations of monetary outcomes, probabilities, as well as their integration.

The arguably most prominent theory to describe human behavior in such situations is *prospect theory* (Kahneman & Tversky, 1979; Tversky & Kahneman, 1992; see Wakker, 2010, for an overview). According to prospect theory, individuals evaluate a decision between lotteries such as A and B by calculating its utilities $U(A)$ and $U(B)$. The core mechanisms that govern that calculation are (a) a reference point relative to which outcomes are evaluated, (b) diminishing sensitivity to larger deviances from the reference point (i.e., the difference between \$10 and \$20 is perceived as larger than the difference between \$1,010 and \$1,020), (c) loss aversion (i.e., losses have a higher impact on utilities relative to gains of the same magnitude), (d) over-weighting of rare events, and (e) under-weighting of probable events. Afterwards, the utilities of A and B are compared with each other and the option with the higher utility is chosen by applying some choice rule.

According to prospect theory, the utility U of a two-outcome mixed lottery $L = \begin{pmatrix} x^+ & x^- \\ p^+ & p^- \end{pmatrix}$ with gain and loss outcomes x^+ and x^- , respectively, and their respective probabilities p^+ and p^- is given by:

$$U(L) = v(x^+)w(p^+) + v(x^-)w(p^-), \quad (4)$$

where $v(\cdot)$ is the (already-mentioned) *value function* and $w(\cdot)$ the *probability-weighting function*. The value function is typically cast as a piecewise power function with parameters α^+ and α^- capturing the diminishing sensitivity in the domains of monetary gains and losses, respectively, and a loss-aversion parameter λ that captures the asymmetries

in the valuation of gains and losses:

$$v(x) = \begin{cases} x^{\alpha^+}, & \text{for } x = x^+, \\ -\lambda|x|^{\alpha^-}, & \text{for } x = x^-, \\ 0, & \text{for } x = 0. \end{cases} \quad (5)$$

Probabilities are assumed to be weighted by an inversely S-shaped function that overweights small probabilities and underweights large probabilities, such as the function proposed by Kahneman and Tversky (1979):

$$w(p) = \frac{p^\gamma}{(p^\gamma + (1-p)^\gamma)^{\frac{1}{\gamma}}}, \quad (6)$$

where γ is the probability-sensitivity parameter. Different parameters γ can also be assigned to the probabilities associated with gains and losses (i.e., γ^+ and γ^-).

Preferences such as A is preferred to B ($A \succ B$) are translated into choice probabilities via a *choice rule* like the logistic choice function:

$$Pr(A) = \frac{1}{1 + e^{-\theta(U(A)-U(B))}}, \quad (7)$$

where the choice-sensitivity parameter θ modulates how differences in utilities affect choice probabilities. Responses are random when $\theta = 0$, and become more deterministic as it increases.

Prospect theory has often been used in the study of individual differences and temporal stability, from risk attitudes to the subjective representation of monetary outcomes and probabilities (e.g., Booij, van Praag, & van de Kuilen, 2009; Broomell & Bhatia, 2014; Kellen, Pachur, & Hertwig, 2016; Scheibehenne & Pachur, 2015). But despite its many merits, prospect theory is sloppy to some degree, and its parameters suffer from well-documented parameter trade-offs, most notably between the outcome-sensitivity parameters α and the choice-sensitivity parameter θ .

Constructing priors for prospect theory

To evaluate the Bayesian variant of the empirical-prior approach in the risky-choice paradigm, we used the data from Walasek and Stewart (2015, Experiment 1a and 1b). In these experiments, participants were faced with a single-lottery accept-reject task, in which they were offered a mixed lottery with two equiprobable outcomes such as $L = \begin{pmatrix} \$20 & -\$12 \\ .50 & .50 \end{pmatrix}$. Participants were asked to decide whether to accept or reject such a lottery, a decision that is assumed to imply a comparison between the lottery and a status quo (with utility 0). This accept-reject task is often used in neuroscientific investigations (e.g., Tom, Fox, Trepel, & Poldrack, 2007; De Martino, Camerer, & Adolphs, 2010; Canessa et al., 2013; Pighin, Bonini, Savadori, Hadjichristidis, & Schena, 2014). Each

participant completed all possible combinations of eight different gains and eight different losses, resulting in a total of 64 trials.

Walasek and Stewart's (2015) study revolved around four different between-subjects conditions that were designed to specifically affect the loss-aversion parameter λ . We will focus on the two conditions that produced the most extreme median λ estimates. In the 40-20 condition ($n = 191$), gain outcomes ranged from \$12 up to \$40 in steps of \$4, whereas losses ranged from \$6 up to \$20 in steps of \$2. The 20-40 condition ($n = 198$) flipped the signs of these outcomes (i.e., gain outcomes became losses and vice-versa). Walasek and Stewart (2015) reported that the λ estimates were generally *above 1* in the 40-20 Condition, indicating loss-averse preferences, and *below 1* in the 20-40 Condition, indicating gain-seeking preferences.

We modeled the data with a streamlined version of prospect theory, in which we assumed that $w(p^+) = w(p^-) = .50$ (see Kellen, Mata, & Davis-Stober, 2017; Levy & Levy, 2002; Quiggin, 1982):

$$P(\text{Accept L}) = \frac{1}{1 + e^{-\frac{\theta}{2}(v(x^+) - v(x^-))}}. \quad (8)$$

Samples from the parameters' posterior distributions were obtained using a No-U-turn sampler (Hoffman & Gelman, 2014) as implemented in *Stan* (Carpenter et al., 2017) via the *RStan* interface (Stan Development Team, 2016b). We ran four randomly initialized chains in parallel for initially 4,000 total iterations, out of which 2,000 were used as a warm-up period to tune the sampler's parameters. These warm-up samples were discarded afterwards. The remaining 2,000 iterations from each chain were thinned and then concatenated, resulting in a total of 1,000 samples. To assess convergence, we used the \hat{R} statistic (Gelman et al., 2013, p. 285) and assumed that convergence was reached if all $\hat{R} \leq 1.01$. If not, we repeated the sampling procedure with twice as many iterations as before, until all parameters converged or a maximum of 64,000 iterations was reached. To avoid singularities in model expectations, we set an upper limit of 2 for parameters α and θ , and of 4 for parameter λ . Also, to avoid numerical over- and underflows, we restricted likelihoods to be between 10^{-7} and $1 - 10^{-7}$. Finally, we used uniform priors that spanned the permitted range of each parameter.

The posterior samples from each individual were then linearly transformed so they would all fall within the $[0, 1]$ range and then inverse-probit transformed into the real space. We used multivariate GMMs (with estimation of the full covariance matrix per multivariate kernel) to

approximate the aggregated individual posterior distribution of the parameters, separately for each of the two conditions. To determine the best-performing GMMs (we considered GMMs with up to ten component distributions), we used leave-one-participant-out cross validation (see Vehtari & Lampinen, 2002, for other variants of cross validation). The parameters of the best-performing GMMs are reported in Table 6 in the Appendix.

The simulation procedure was similar to the one used in the first part of the paper, but extended by one additional factor: For each of the two conditions, we generated data from a uniform distribution within the restricted parameter boundaries and from each of the two fitted prior distributions. Afterwards, we obtained samples from the posterior distributions using a uniform prior, the prior obtained from fitting the 40-20 condition, or the prior obtained from fitting the 20-40 condition. These samples were obtained using a differential-evolution sampler (Ter Braak & Vrugt, 2008) as implemented in *BAYESIANTOOLS* package (Hartig, Minunno, & Paul, 2017). Consequently, this resulted in a 2 (condition) \times 3 (ground-truth prior distribution) \times 3 (used prior) simulation design. Within each cell, we obtained a total of 1,000 observations.

As dependent variables, we used the coverage ratio, the r^2 across individuals of the true parameters and the mode of the respective posterior distributions, and the proportion of true parameters included in the 95% credible interval, $P(95\%CI)$. A low coverage ratio and a proportion close to .95 of parameters included in the 95% credible interval hint towards good parameter identifiability, and a high r^2 reflects a good recovery of the rank ordering of parameters across individuals.

Results

Empirical prior in the 40-20 condition

The empirical distributions of parameters in the 40-20 condition mostly follow the expectations about prospect theory parameter distributions reported in the literature (Booij et al., 2009). We observed a slight tendency towards risk aversion (i.e., the posterior distribution of α peaks slightly below 1), a tendency towards loss aversion (i.e., most of the posterior mass of λ is above 1), and choices that are stochastic (i.e., the peak of the posterior distribution of θ tends towards 0). It is noteworthy that the distribution of λ is multimodal. It looks like there are at least two peaks, one around a value slightly below 2 and one at 1 (i.e., loss neutrality). See Fig. 5 (main diagonal, gray line) for fine-grained histograms of the marginal distributions of the parameters.




Fig. 5 Posterior distribution across individuals in the 40-20 condition of the risky choice paradigm and the *Gaussian mixture model* (GMM) that was fitted to it. The main diagonal depicts the marginal distributions of the parameters of the posterior samples (*gray line*) and the fitted mixture model (*black line*). Lower diagonal elements show the pair plots of the posterior samples, upper diagonal elements the pair plots of the fitted mixture model. Parameters α , λ , and θ reflect risk aversion, loss aversion, and the scaling parameter of the logistic choice function, respectively

The inspection of joint parameter distributions (see Fig. 5, lower diagonal elements) reveals strong dependencies. The negative, curvilinear dependency between α and θ resembles the dependency reported by Scheibehenne and Pachur (2015). The multimodality of the λ parameter makes it difficult to interpret its interdependencies. Disregarding the peak of λ at 1, at which the parameter has no influence on decisions (and thus should not correlate with any other parameter), λ seems to be positively correlated with α and negatively correlated with θ , a pattern that is not very surprising: Larger values of λ lead to a larger influence of losses on the decision variable, which can be partially compensated for by also increasing the symmetrical scaling of both losses and gains (α). These large values in the decision variable, in turn, would lead to more deterministic choices, which can be scaled down with lower values of θ .

The best-fitting GMM in the 40-20 condition turned out to be a mixture of three components (see Fig. 5, main

diagonal, black line for the marginal distributions of the best-fitting GMM and Table 4 in the Appendix for the BICs for all numbers of mixtures). Whereas the distributions of α and θ were approximated very closely, the multimodality of λ cannot be well accommodated with this solution.⁸ Except for the fan-like correlation of λ with α , the GMM was able to closely approximate the covariations found among other parameter pairings.

Empirical prior in the 20–40 condition

The empirical distributions of parameters in the 20-40 condition reflected the experimental manipulations reported by Walasek and Stewart (2015). We observed risk neutrality (i.e., the posterior distribution of α peaks at around 1), a

⁸As a robustness check, we fit up to 30 mixture components to the data of condition 40-20, but the fit of the mixture components got worse with each added component.




Fig. 6 Posterior distribution across individuals in the 20-40 condition of the risky choice paradigm and the Gaussian mixture model that was fitted to it. The main diagonal depicts the marginal distributions of the parameters of the posterior samples (*gray line*) and the fitted mixture model (*black line*). Lower diagonal elements show the pair plots of the posterior samples, upper diagonal elements the pair plots of the fitted mixture model. Parameters α , λ , and θ reflect risk aversion, loss aversion, and the scaling parameter of the logistic choice function, respectively

slight tendency towards gain seeking (i.e., the distribution of λ peaks slightly below 1), and stochasticity of choices. All distributions were found to be unimodal, making it easier for the GMMs to approximate them. See Fig. 6 (main diagonal, gray line) for fine-grained histograms of the marginal distributions of the parameters. The inspection of the parameter distributions (see Fig. 6, lower diagonal elements) only revealed a strong dependency between α and θ . This ridge-like relationship is very similar to the one observed in the distributions of the 40-20 conditions. Otherwise, there were almost no dependencies observable. Note that this lack of dependencies mainly results from the fact that neither an α of 1 nor a λ of 1 have any influence on the decision variable. Therefore, distributions that are strongly peaked around 1 cannot sensibly covary with other parameters.

The best-fitting GMM to the posterior distribution of parameters in the 20-40 condition was a mixture of four Gaussians (see Fig. 6, main diagonal, black line for the marginal distributions of the best-fitting GMM and Table 4 for the BICs for all numbers of mixtures). Apart from the height of the peak of λ , all other aspects of the

empirical distributions, including the covariations, were well approximated.

Simulation results

We simulated 1,000 virtual participants from each of 2 (experimental condition) \times 3 (ground-truth prior) = 6 factor combinations. We then refitted the data coming from each of these virtual participants under three different conditions: (a) using a uniform prior, (b) the empirical prior obtained for the 40-20 condition, and (c) the empirical prior obtained for the 20-40 condition. Just as in the first part of the paper, we first report global results aggregated across all factors, only then turning to the effects of matching conditions and priors, and the influence of (mis)matches between them.

Results reported in Table 3 show that across all experimental conditions, population distributions, and priors, parameter λ was recovered best, followed by α and θ . This rank order holds for both the coverage ratio and r^2 . However, these results were far from satisfactory, as $r^2(\lambda) = .30$, $r^2(\alpha) = .11$, and $r^2(\theta) = .02$ across conditions. Also, the

Table 3 Parameter recoverability of the fully Bayesian empirical-prior method using streamlined prospect theory

Ground-truth distribution & condition	Used prior	Parameter	$P(95\%CI)$	Coverage ratio [M, Md, SD]	r^2
40-20	40-20	α	.59	[.20, .20, .04]	.42
		λ	.66	[.16, .15, .06]	.74
		θ	.57	[.50, .49, .07]	.07
	20-40	α	.44	[.17, .16, .05]	.33
		λ	.31	[.10, .10, .02]	.72
		θ	.46	[.45, .44, .06]	.02
	Uniform	α	.78	[.49, .48, .22]	.29
		λ	.78	[.36, .36, .22]	.52
		θ	.85	[.78, .95, .33]	.06
20-40	40-20	α	.57	[.23, .23, .02]	.43
		λ	.76	[.12, .09, .08]	.69
		θ	.52	[.52, .50, .07]	.00
	20-40	α	.57	[.23, .24, .02]	.43
		λ	.68	[.07, .07, .02]	.59
		θ	.48	[.48, .46, .06]	.00
	Uniform	α	.90	[.63, .61, .22]	.19
		λ	.87	[.32, .13, .31]	.51
		θ	.93	[.91, .98, .21]	.01

$P(95\%CI)$ = Proportion of true parameters included in the 95% credible interval. Coverage ratio = range of the 95% credible interval divided by the parameter support (i.e., upper boundary – lower boundary). r^2 = shared variance across individuals of the true parameters and the mode of the respective posterior distributions. α = curvature of the utility function. λ = parameter governing the importance of losses relative to gains. θ = scaling parameter of the logistic choice function

95% credible intervals included the true parameter values at much lower rates than expected.

In cases where the prior used matched the data-generating population and the condition, parameter recovery was on average slightly better. In the case of the 40–20 condition, this led to a significantly lower coverage ratio for both λ ($M = .16, Md = .15, SD = .06$) and α ($M = .20, Md = .20, SD = .04$). While the coverage ratio for θ decreased as well, it remained somewhat unsatisfactory as it still spanned roughly half of the range of possible values ($M = .50, Md = .49, SD = .07$). The pattern of rank stability shows a somewhat different picture though: Whereas the correlation between the ground-truth values and the posterior modes of λ ($r^2 = .75$) and α ($r^2 = .42$) improved dramatically (compared to the aggregated values), it became worse for θ ($r^2 = .07$). The proportion of

ground-truth parameters included in the 95% credible interval barely changed ($M_{min} = .57, M_{max} = .66$). Very similar results were found in the 20-40 condition.

The scenario in which the prior matched the data-generating population as well as the condition is the most optimistic one. This notion is important when considering the rates at which the 95% credible intervals included the true parameter values. With increasing sloppiness of a model and fully Bayesian parameter estimation, one would expect the credible intervals to widen without affecting that rate. However, as shown in Fig. 7, the credible intervals missed the true values much more often than they should. This result shows that the challenges created by parameter non-identifiability and sloppiness are not automatically addressed by a fully Bayesian treatment. Crucially, the posterior distributions appear well-behaved and show no signs that anything might be wrong with the model specification (see Fig. 7).

Let us now turn to the (perhaps more realistic) cases in which there was a mismatch between the ground truth and the modeling assumptions. Here, we report the results from mixed mismatching between condition and prior used (i.e., participants stem from the ground-truth prior from the condition from which the priors were obtained, while the prior that is used during re-fitting varies). Table 3 reports all dependent variables for all the combinations analyzed here. When using the empirical prior from 20-40 condition in the fitting of data from the 20-40 condition, we observe lower coverage ratios together with lower proportions of true values included in the 95% credible interval. Both variables improved when the uniform prior was used instead. Results were somewhat similar when the mismatching data came from the 40-20 condition.

General discussion

The present work evaluated the empirical-prior approach for obtaining informative priors, which has been proposed as a way to deal with problems concerning parameter non-identifiability and model sloppiness (Gershman, 2016). Using the reinforcement-learning data originally reported by Gershman (2015), we first tested how the point-estimate variant of the empirical-prior approach fared in comparison with simple MLE. We found that neither approach provided satisfactory results and that neither one of them consistently outperformed the other. We then considered potential variations of Gershman’s experimental design as ways to improve recoverability. Modest but encouraging improvements were observed when increasing the number of trials per block or the number of options made available to the participants. To assess whether the rather




Fig. 7 Violin plots for 20 individual posterior distributions of random virtual participants in the risky choice paradigm. The ground-truth parameters stemmed from the 40–20 prior in the 40–20 condition and were fitted using the 40–20 empirical prior. *Thick bars* represent the 95% credible interval. *Vertical red line* represents the ground-truth parameter. Note that in expectation, the 95% credible interval should include the red line in .95 of the cases (i.e., only one out of 20 red lines for each parameter should lie outside the credible interval). In that sample, eight, five, and ten ground truths of α , λ , and θ , respectively, lie outside the 95% credible interval. Parameters α , λ , and θ reflect risk aversion, loss aversion, and the scaling parameter of the logistic choice function, respectively. The *violin plots* have been scaled vertically to fit their respective cells

poor performance of the point-estimate empirical-prior method was specific to its application to reinforcement-learning models (and the reliance on point estimates), we developed a fully Bayesian extension to the method and tested it in a streamlined variant of prospect theory (Kahneman & Tversky, 1979). In line with the results we obtained so far, we again did not observe a general advantage of the empirical-prior method. Important, we found that the true parameter values were often missed by the estimates' respective 95% credible intervals, even when under a best-case scenario in which both the model and priors are “true”. This result goes counter to the expectation that parameter non-identifiability and sloppiness issues are well captured by the posterior distributions such that they should simply lead to wider posteriors. Instead, we

often find posterior distributions that are concentrated in regions that do not include the true data-generating values.

Fitting data from an experiment and plugging the resulting parameter distributions as informative priors into a separate model-fitting procedure is an elegant and easy-to-implement procedure. Unfortunately, the informativeness of these informative priors is limited, and the method does not help solve the problem it was designed for. We showed that even when the priors used for the model-fitting procedure (be it using MAP or fully Bayesian estimation) are aligned with the true underlying parameter distributions, there are no systematic advantages of using informative over uniform priors. In case of a mismatch, which in empirical settings is likely to be the case, the ability to recover parameters

can drop dramatically, even for rather simple models. Given that the true underlying parameter distributions cannot be recovered to a satisfactory degree, the ability to compare group-level differences is also compromised.

In the case of Gershman's (2015) baseline experiment, we found that the main culprit for the poor recoverability was the limited informativeness of the data. On average, the parameters posterior distributions were well dispersed across the ranges of possible values, making it practically impossible to reasonably interpret any point estimates obtained through model fitting. Although some of these problems could be ameliorated by extending the experimental design, such extensions can also introduce their own set of practical problems. For instance, the increase of the number of options implies an increase in terms of tasks demands, which can in turn lead to individual preference profiles that models have trouble accounting for (e.g., Steingrover et al., 2014, for a demonstration in the Iowa gambling task; Bechara et al., 1994). Similarly, full feedback can lead to behavioral phenomena that are either unique to such scenarios, like attention allocation to foregone outcomes (e.g., Ashby & Rakow, 2016), or that at least differ considerably from what is found in the case of partial feedback (e.g., Plonsky, Teodorescu, & Erev, 2015, Plonsky & Erev, 2017; Yechiam, Stout, Busemeyer, Rock, & Finn, 2005).

Inversion versus inference

The concepts of parameter identifiability, recoverability, and model sloppiness discussed here are instrumental when attempting to infer the ground truth from data generated by it. Within the context of this *question of inversion*, identifiability and recoverability are of utmost importance, as without them it is impossible to draw correct conclusion about the underlying cognitive processes. For example, the β and θ parameters of the evaluated reinforcement-learning models and prospect theory, respectively, had the lowest recoverability of rank orders as reflected in r^2 values close to 0. In light of such poor parameter recovery, a relatively high estimated parameter value was not predictive of whether the "true" choice consistency of the respective virtual participant was high or low.

However, such concerns do not carry over wholesale when, for instance, one frames the problem of parameter estimation as a *question of inference*. In this context, the ability to recover the ground truth cedes the center stage to the coherence of our relative support for the different hypotheses. For example, consider a model-selection scenario in which data are generated from a complex model, but turn out to still be somewhat likely

under a much simpler candidate model. A greater support for the simpler model is a sensible conclusion here as this is the model that provides the best trade-off between goodness of fit and parsimony, even if it did not generate the data (Lee, *forthcoming*, pp. 60–61).⁹ After all, models that are "wrong" (e.g., simpler than the generative model) can still be useful in predicting behavior (e.g., Lee & Webb, 2005). Nevertheless, it would be unwise to assume that even under such framing, we can completely divorce ourselves from any concerns related with identifiability and sloppiness. After all, it is still sensible to carefully evaluate the roles that the different parameters in a model can play, and how these can be ascertained under different experimental designs (Broomell & Bhatia, 2014). And even if one is ultimately not attempting to recover some ground truth, parameter-recovery exercises in which a ground truth is known can be seen as a sandbox that helps us to understand current difficulties in disentangling the roles parameters play, and develop ways to overcome them.

Conclusions

Computational models are popular tools to develop and test psychological theories of cognition. For them to also be *useful* tools, it is important to ensure that the parameters obtained from fitting models to data provide an accurate characterization of the underlying cognitive processes. If the data are not suited to inform us about the model parameters, as in the simple probability learning task used by Gershman (2015), then this requirement is not fulfilled. Informative priors used during the model fitting procedure can be helpful for estimating parameters (see Lee & Vanpaemel, 2017), however, they do not constitute a panacea for the identifiability or sloppiness problems that often arise when

⁹The concerns with model identifiability are also minor when framing the problem of parameter estimation as a mechanism for obtaining model expectations. For example, the very same reinforcement-learning models discussed here are often applied to obtain model-based estimates (e.g., reward expectations at the end of a learning phase) that can be compared with different variables, such as the blood-oxygenation-level dependent activity in the brain (e.g., Leong, Radulescu, Daniel, DeWoskin, & Niv, 2017; Jocham, Klein, & Ullsperger, 2011; Frank et al., 2015; Niv, Edlund, Dayan, & O'Doherty, 2012; see Lee, Seo, & Jung, 2012, for an overview). These results are entirely unaffected by non-identifiability (and only weakly affected by sloppiness), as all of the infinite parameter combinations that result in equal likelihoods necessarily stem from identical reward expectations (up to a scaling factor that is proportional to the scaling parameter of the error model).

using non-informative experimental designs. In contrast, simple adjustments in the experimental design can often improve parameter recoverability. Based on these results, we conclude that researchers should invest more of their efforts in assessing and improving the information content of their experimental designs instead of relying on statistical methods after the fact. In the end, whether empirical priors help or not (and how much) is a shot in the dark, as they

only seem to help in some of the rare cases in which there is a match between the priors and population distribution.

Author Note This research was supported by the Swiss National Science Foundation (SNSF Grant 100014_165591 to David Kellen). We thank Henrik Singmann, Gregory E. Cox, Samuel J. Gershman, and Michael D. Lee for valuable comments on a previous version of this manuscript. We also thank Samuel J. Gershman and Lukasz Walasek for making their data available. Simulation results and analysis scripts are available on the Open Science Framework (<https://osf.io/2ws78/>).

Appendix

Table 4 Bayesian information criteria of the Gaussian mixture models for each of the substantive models

Model	Parameter/Condition	Number of mixture kernels									
		1	2	3	4	5	6	7	8	9	10
\mathcal{M}_1	η	907	876	871	731	739	746	754	762	770	781
	β	880	689	607	591	600	608	602	611	623	633
\mathcal{M}_2	η^+	891	851	848	858	867	872	881	897	912	937
	η^-	900	849	694	713	695	714	725	735	747	752
	β	932	808	706	709	720	698	699	704	717	722
\mathcal{M}_3	η	861	805	800	801	813	817	826	828	838	851
	β	729	536	534	542	547	560	569	581	594	602
	ω	804	364	265	222	200	197	206	215	226	232
\mathcal{M}_4	η^+	910	847	812	790	792	798	811	817	820	821
	η^-	883	823	768	755	757	763	779	789	795	806
	β	728	581	585	592	600	607	615	622	633	642
	ω	794	294	156	82	72	55	49	52	59	66
PT	40-20	130	120	119	120	120	122	123	124	125	127
	20-40	127	113	114	110	113	112	114	115	116	117

See “Models” and “Constructing priors for prospect theory” for model specifications of the reinforcement-learning models and prospect theory (PT), respectively. The Bayesian information criteria were obtained by fitting the Gaussian mixture models to the marginal distributions of each of the parameters 500 times (reinforcement-learning models) or by using leave-one-participant-out cross validation (prospect theory). Highlighted cells are the best-fitting solutions for the respective parameters. The reported values for prospect theory are scaled down and rounded to 10^4 .

Table 5 Mixture components of the empirical priors for each of the reinforcement-learning models

Model	Parameter	Mixture component	μ	σ	w	
\mathcal{M}_1	η	1	-0.13	0.76	.53	
		2	6.22	0.47	.22	
		3	-2.53	0.40	.14	
		4	-5.62	0.40	.10	
	β	1	-1.73	0.52	.71	
		2	6.14	0.65	.14	
		3	-6.10	0.31	.13	
		4	0.60	0.49	.03	
	\mathcal{M}_2	η^+	1	0.83	4.79	.44
			2	-0.66	0.97	.37
			3	-2.35	0.33	.19
		η^-	1	-0.14	1.39	.71
2			6.17	0.30	.21	
3			-6.34	0.08	.08	
β		1	-1.55	0.44	.49	
		2	6.00	0.51	.25	
		3	-5.96	0.29	.10	
		4	-2.55	0.29	.07	
		5	2.33	2.10	.06	
		6	2.95	0.17	.03	
\mathcal{M}_3	η	1	-0.46	1.11	.67	
		2	3.52	3.71	.30	
		3	-6.20	0.17	.04	
	β	1	-1.82	0.61	.81	
		2	0.29	2.04	.13	
		3	5.95	0.45	.06	
	ω	1	-0.10	0.34	.36	
		2	-0.02	0.04	.24	
		3	0.00	0.00	.18	
		4	-6.25	0.51	.12	
		5	6.20	0.38	.06	
		6	-1.63	0.32	.04	
\mathcal{M}_4	η^+	1	-0.77	1.29	.64	
		2	5.69	0.80	.24	
		3	-5.41	0.44	.07	
		4	-6.35	0.02	.05	
	η^-	1	0.10	0.81	.62	
		2	-5.04	1.12	.15	
		3	6.20	0.23	.15	
		4	4.09	0.96	.08	
	β	1	-1.92	0.54	.71	
		2	1.68	2.29	.29	

Table 5 (continued)

Model	Parameter	Mixture component	μ	σ	w
\mathcal{M}_1	ω	1	-0.09	0.19	.33
		2	-0.02	0.03	.20
		3	0.00	0.00	.16
		4	-0.96	2.52	.11
		5	-6.32	0.08	.10
		6	6.00	0.46	.05
		7	0.00	0.00	.04

See “Models” for model specifications. The mixtures are defined on the real space; For the correct parameter space, the mixture draws have to be transformed according to the formulae provided in “Empirical priors”. Note that mixture components 3 and 7 for the ω parameter of model \mathcal{M}_4 are, in fact, not identical, but only differ after the second decimal. w = mixture weight

Table 6 Mixture components of the empirical priors for prospect theory

Condition	Parameter(s)	Mixture 1	Mixture 2	Mixture 3	Mixture 4	
40-20	μ	α	-0.17	0.17	0.13	
		λ	-0.19	0.15	-0.25	
		θ	0.37	-0.82	0.05	
	σ_X^2	α	0.05	0.15	0.91	
		λ	0.14	0.49	1.12	
		θ	0.61	0.39	1.20	
$\sigma_{X,Y}$	α, λ	-0.01	-0.02	0.02		
	α, θ	-0.09	-0.20	-0.28		
	λ, θ	-0.03	-0.07	0.01		
20-40	μ	w	.39	.31	.29	
		α	-0.01	-0.19	0.52	0.78
		λ	-0.80	-0.24	-0.64	-0.70
	σ_X^2	θ	0.14	-0.03	-1.09	0.46
		α	0.08	0.68	0.28	0.48
		λ	0.06	1.29	0.16	0.02
$\sigma_{X,Y}$	θ	0.59	1.05	0.30	1.05	
	α, λ	-0.02	0.18	-0.08	0.00	
	α, θ	-0.13	-0.19	-0.20	-0.45	
w	λ, θ	0.01	-0.18	-0.01	-0.03	
	w	.49	.19	.18	.15	

The mixtures are defined on the real space; For the correct parameter space, the mixture draws have to be inverse-probit transformed and then multiplied with the parameter boundaries of 2, 4, and 2 for α , λ , and θ , respectively. w = mixture weight

References

- Ahn, W.-Y., Krawitz, A., Kim, W., Busemeyer, J. R., & Brown, J. W. (2011). A model-based fMRI analysis with hierarchical Bayesian parameter estimation. *Journal of Neuroscience, Psychology, and Economics*, 4, 95–110. <https://doi.org/10.1037/a0020684>
- Ahn, W.-Y., Vasilev, G., Lee, S.-H., Busemeyer, J. R., Kruschke, J. K., Bechara, A., & Vassileva, J. (2014). Decision-making in stimulant and opiate addicts in protracted abstinence: Evidence from computational modeling with pure users. *Frontiers in Psychology*, 5, 1–15. <https://doi.org/10.3389/fpsyg.2014.00849>
- Ashby, N. J. S., & Rakow, T. (2016). Eyes on the prize? Evidence of diminishing attention to experienced and foregone outcomes in repeated experiential choice. *Journal of Behavioral Decision Making*, 29, 183–193. <https://doi.org/10.1002/bdm.1872>
- Bamber, D., & van Santen, J. P. (1985). How many parameters can a model have and still be testable? *Journal of Mathematical Psychology*, 29, 443–473. [https://doi.org/10.1016/0022-2496\(85\)90005-7](https://doi.org/10.1016/0022-2496(85)90005-7)
- Barron, G., & Erev, I. (2003). Small feedback-based decisions and their limited correspondence to description-based decisions. *Journal of Behavioral Decision Making*, 16, 215–233. <https://doi.org/10.1002/bdm.443>
- Batchelder, W. H., & Riefer, D. M. (1990). Multinomial processing models of source monitoring. *Psychological Review*, 97, 548–564. <https://doi.org/10.1037/0033-295X.97.4.548>
- Bechara, A., Damasio, A. R., Damasio, H., & Anderson, S. W. (1994). Insensitivity to future consequences following damage to human prefrontal cortex. *Cognition*, 50, 7–15. [https://doi.org/10.1016/0010-0277\(94\)90018-3](https://doi.org/10.1016/0010-0277(94)90018-3)
- Booij, A. S., van Praag, B. M. S., & van de Kuilen, G. (2009). A parametric analysis of prospect theory's functionals for the general population. *Theory and Decision*, 68, 115–148. <https://doi.org/10.1007/s11238-009-9144-4>
- Broomell, S. B., & Bhatia, S. (2014). Parameter recovery for decision modeling using choice data. *Decision*, 1, 252–274. <https://doi.org/10.1037/dec0000020>
- Brown, K. S., & Sethna, J. P. (2003). Statistical mechanical approaches to models with many poorly known parameters. *Physical Review E*, 68, 021904. <https://doi.org/10.1103/PhysRevE.68.021904>
- Buchner, A., & Erdfelder, E. (2005). Word frequency of irrelevant speech distractors affects serial recall. *Memory & Cognition*, 33, 86–97. <https://doi.org/10.3758/BF03195299>
- Canessa, N., Crespi, C., Motterlini, M., Baud-Bovy, G., Chierchia, G., Pantaleo, G., & Cappa, S. F. (2013). The functional and structural neural basis of individual differences in loss aversion. *Journal of Neuroscience*, 33, 14307–14317. <https://doi.org/10.1523/JNEUROSCI.0497-13.2013>
- Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., & Riddell, A. (2017). Stan: A probabilistic programming language. *Journal of Statistical Software*, 76, 1–32. <https://doi.org/10.18637/jss.v076.i01>
- Chase, H. W., Kumar, P., Eickhoff, S. B., & Dombrovski, A. Y. (2015). Reinforcement learning models and their neural correlates: An activation likelihood estimation meta-analysis. *Cognitive, Affective, & Behavioral Neuroscience*. <https://doi.org/10.3758/s13415-015-0338-7>
- Cousineau, D., & Hélie, S. (2013). Improving maximum likelihood estimation using prior probabilities: A tutorial on maximum a posteriori estimation and an examination of the Weibull distribution. *Tutorials in Quantitative Methods for Psychology*, 9, 61–71. <https://doi.org/10.20982/tqmp.09.2.p061>
- Dayan, P., & Balleine, B. W. (2002). Reward, motivation, and reinforcement learning. *Neuron*, 36, 285–298. [https://doi.org/10.1016/S0896-6273\(02\)00963-7](https://doi.org/10.1016/S0896-6273(02)00963-7)
- Dayan, P., & Daw, N. D. (2008). Decision theory, reinforcement learning, and the brain. *Cognitive, Affective, & Behavioral Neuroscience*, 8, 429–453. <https://doi.org/10.3758/CABN.8.4.429>
- De Martino, B., Camerer, C. F., & Adolphs, R. (2010). Amygdala damage eliminates monetary loss aversion. *Proceedings of the National Academy of Sciences*, 107, 3788–3792. <https://doi.org/10.1073/pnas.0910230107>
- Erev, I., & Barron, G. (2005). On adaptation, maximization, and reinforcement learning among cognitive strategies. *Psychological Review*, 112, 912–931. <https://doi.org/10.1037/0033-295X.112.4.912>
- Frank, M. J., Gagne, C., Nyhus, E., Masters, S., Wiecki, T. V., Cavanagh, J. F., & Badre, D. (2015). fMRI and EEG predictors of dynamic decision parameters during human reinforcement learning. *Journal of Neuroscience*, 35, 485–494. <https://doi.org/10.1523/JNEUROSCI.2036-14.2015>
- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A., & Rubin, D. B. (2013). *Bayesian data analysis*, (3rd ed.). Boca Raton: CRC Press.
- Gershman, S. J. (2015). Do learning rates adapt to the distribution of rewards? *Psychonomic Bulletin & Review*, 22, 1320–1327. <https://doi.org/10.3758/s13423-014-0790-3>
- Gershman, S. J. (2016). Empirical priors for reinforcement learning models. *Journal of Mathematical Psychology*, 71, 1–6. <https://doi.org/10.1016/j.jmp.2016.01.006>
- Hartig, F., Minunno, F., & Paul, S. (2017). *BayesianTools: General-purpose MCMC and SMC samplers and tools for Bayesian statistics*. R package version 0.1.3. Retrieved from https://github.com/florianhartig/bayesian_tools.
- Hoffman, M. D., & Gelman, A. (2014). The No-U-Turn sampler: Adaptively setting path lengths in Hamiltonian Monte Carlo. *Journal of Machine Learning Research*, 15, 1593–1623. arXiv:1111.4246.
- Hulme, C., Roodenrys, S., Schweickert, R., Brown, G. D. A., et al. (1997). Word-frequency effects on short-term memory tasks: Evidence for a reintegration process in immediate serial recall. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 23, 1217–1232. <https://doi.org/10.1037/0278-7393.23.5.1217>
- Humphries, M. A., Bruno, R., Karpievitch, Y., & Wotherspoon, S. (2015). The expectancy valence model of the Iowa gambling task: Can it produce reliable estimates for individuals? *Journal of Mathematical Psychology*, 64–65, 17–34. <https://doi.org/10.1016/j.jmp.2014.10.002>
- Jocham, G., Klein, T. A., & Ullsperger, M. (2011). Dopamine-mediated reinforcement learning signals in the striatum and ventromedial prefrontal cortex underlie value-based choices. *Journal of Neuroscience*, 31, 1606–1613. <https://doi.org/10.1523/JNEUROSCI.3904-10.2011>
- Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47, 263–292. <https://doi.org/10.2307/1914185>
- Katahira, K. (2016). How hierarchical models improve point estimates of model parameters at the individual level. *Journal of Mathematical Psychology*, 73, 37–58. <https://doi.org/10.1016/j.jmp.2016.03.007>
- Kellen, D., Mata, R., & Davis-Stober, C. P. (2017). Individual classification of strong risk attitudes: An application across lottery types and age groups. *Psychonomic Bulletin & Review*, 24, 1341–1349. <https://doi.org/10.3758/s13423-016-1212-5>

- Kellen, D., Pachur, T., & Hertwig, R. (2016). How (in)variant are subjective representations of described and experienced risk and rewards? *Cognition*, *157*, 126–138. <https://doi.org/10.1016/j.cognition.2016.08.020>
- Lee, D., Seo, H., & Jung, M. W. (2012). Neural basis of reinforcement learning and decision making. *Annual Review of Neuroscience*, *35*, 287–308. <https://doi.org/10.1146/annurev-neuro-062111-150512>
- Lee, M. D. (forthcoming). Bayesian methods in cognitive modeling. In Wixted, J. T. (Ed.) *The Stevens' handbook of experimental psychology and cognitive neuroscience* (4th edition, volume 5: Methodology). New York: Wiley.
- Lee, M. D., & Vanpaemel, W. (2017). Determining informative priors for cognitive models. *Psychonomic Bulletin & Review*. <https://doi.org/10.3758/s13423-017-1238-3>.
- Lee, M. D., & Webb, M. R. (2005). Modeling individual differences in cognition. *Psychonomic Bulletin & Review*, *12*, 605–621. <https://doi.org/10.3758/BF03196751>
- Leong, Y. C., Radulescu, A., Daniel, R., DeWoskin, V., & Niv, Y. (2017). Dynamic interaction between reinforcement learning and attention in multidimensional environments. *Neuron*, *93*, 451–463. <https://doi.org/10.1016/j.neuron.2016.12.040>
- Levy, H., & Levy, M. (2002). Experimental test of the prospect theory value function: A stochastic dominance approach. *Organizational Behavior and Human Decision Processes*, *89*, 1058–1081. [https://doi.org/10.1016/S0749-5978\(02\)00011-0](https://doi.org/10.1016/S0749-5978(02)00011-0)
- Lewandowsky, S., & Farrell, S. (2010). *Computational modeling in cognition: Principles and practice*. Thousand Oaks: Sage Publications Inc.
- Li, S.-C., Lewandowsky, S., & DeBrunner, V. E. (1996). Using parameter sensitivity and interdependence to predict model scope and falsifiability. *Journal of Experimental Psychology: General*, *125*, 360–369. <https://doi.org/10.1037/0096-3445.125.4.360>
- Moran, R. (2016). Thou shalt identify! The identifiability of two high-threshold models in confidence-rating recognition (and super-recognition) paradigms. *Journal of Mathematical Psychology*, *73*, 1–11. <https://doi.org/10.1016/j.jmp.2016.03.002>
- Mullen, K., Ardia, D., Gil, D., Windover, D., & Cline, J. (2011). DEoptim : An R package for global optimization by differential evolution. *Journal of Statistical Software*, *40*, 1–17. <https://doi.org/10.18637/jss.v040.i06>
- Nilsson, H., Rieskamp, J., & Wagenmakers, E.-J. (2011). Hierarchical Bayesian parameter estimation for cumulative prospect theory. *Journal of Mathematical Psychology*, *55*, 84–93. <https://doi.org/10.1016/j.jmp.2010.08.006>
- Niv, Y., Daniel, R., Geana, A., Gershman, S. J., Leong, Y. C., Radulescu, A., & Wilson, R. C. (2015). Reinforcement learning in multidimensional environments relies on attention mechanisms. *Journal of Neuroscience*, *35*, 8145–8157. <https://doi.org/10.1523/JNEUROSCI.2978-14.2015>
- Niv, Y., Edlund, J. A., Dayan, P., & O'Doherty, J. P. (2012). Neural prediction errors reveal a risk-sensitive reinforcement-learning process in the human brain. *Journal of Neuroscience*, *32*, 551–562. <https://doi.org/10.1523/JNEUROSCI.5498-10.2012>
- Pighin, S., Bonini, N., Savadori, L., Hadjichristidis, C., & Schena, F. (2014). Loss aversion and hypoxia: Less loss aversion in oxygen-depleted environment. *Stress*, *17*, 204–210. <https://doi.org/10.3109/10253890.2014.891103>
- Plonsky, O., & Erev, I. (2017). Learning in settings with partial feedback and the wavy recency effect of rare events. *Cognitive Psychology*, *93*, 18–43. <https://doi.org/10.1016/j.cogpsych.2017.01.002>
- Plonsky, O., Teodorescu, K., & Erev, I. (2015). Reliance on small samples, the wavy recency effect, and similarity-based learning. *Psychological Review*, *122*, 621–647. <https://doi.org/10.1037/a0039413>
- Quiggin, J. (1982). A theory of anticipated utility. *Journal of Economic Behavior & Organization*, *3*, 323–343. arXiv:1011.1669v3. [https://doi.org/10.1016/0167-2681\(82\)90008-7](https://doi.org/10.1016/0167-2681(82)90008-7)
- R Development Core Team (2008). R: A language and environment for statistical computing. R Foundation for Statistical computing, Vienna, Austria. Retrieved from <http://www.r-project.org>.
- Scheibehenne, B., & Pachur, T. (2015). Using Bayesian hierarchical parameter estimation to assess the generalizability of cognitive models of choice. *Psychonomic Bulletin & Review*, *22*, 391–407. <https://doi.org/10.3758/s13423-014-0684-4>
- Schmittmann, V. D., Dolan, C. V., Raijmakers, M. E., & Batchelder, W. H. (2010). Parameter identification in multinomial processing tree models. *Behavior Research Methods*, *42*, 836–846. <https://doi.org/10.3758/BRM.42.3.836>
- Schulze, C., van Ravenzwaaij, D., & Newell, B. R. (2015). Of matchers and maximizers: How competition shapes choice under risk and uncertainty. *Cognitive Psychology*, *78*, 78–98. <https://doi.org/10.1016/j.cogpsych.2015.03.002>
- Schwarz, G. (1978). Estimating the dimension of a model. *The Annals of Statistics*, *6*, 461–464. <https://doi.org/10.1214/aos/1176344136>
- Schweickert, R. (1993). A multinomial processing tree model for degradation and redintegration in immediate recall. *Memory & Cognition*, *21*, 168–175. <https://doi.org/10.3758/BF03202729>
- Stan Development Team (2016a). PyStan: The Python interface to Stan. Retrieved from <http://mc-stan.org>.
- Stan Development Team (2016b). RStan: The R interface to Stan. Retrieved from <http://mc-stan.org>.
- Steingrover, H., Wetzels, R., & Wagenmakers, E.-J. (2013). Validating the PVL-Delta model for the Iowa gambling task. *Frontiers in Psychology*, *4*, 1–17. <https://doi.org/10.3389/fpsyg.2013.00898>
- Steingrover, H., Wetzels, R., & Wagenmakers, E.-J. (2014). Absolute performance of reinforcement-learning models for the Iowa gambling task. *Decision*, *1*, 161–183. <https://doi.org/10.1037/dec0000005>
- Sutton, R. S., & Barto, A. G. (1998). *Reinforcement learning: An introduction*. Cambridge: MIT Press.
- Ter Braak, C. J. F., & Vrugt, J. A. (2008). Differential evolution Markov chain with snooker updater and fewer chains. *Statistics and Computing*, *18*, 435–446. <https://doi.org/10.1007/s11222-008-9104-9>
- Tom, S. M., Fox, C. R., Trepel, C., & Poldrack, R. A. (2007). The neural basis of loss aversion in decision-making under risk. *Science*, *315*, 515–518. <https://doi.org/10.1126/science.1134239>
- Tversky, A., & Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty*, *5*, 297–323. <https://doi.org/10.1007/BF00122574>
- Vehtari, A., & Lampinen, J. (2002). Bayesian model assessment and comparison using cross-validation predictive densities. *Neural Computation*, *14*, 2439–2468. <https://doi.org/10.1162/08997660293292>
- Wakker, P. P. (2010). *Prospect theory: For risk and ambiguity*. Cambridge: Cambridge University Press.
- Walasek, L., & Stewart, N. (2015). How to make loss aversion disappear and reverse: Tests of the decision by sampling origin of loss aversion. *Journal of Experimental Psychology: General*, *144*, 7–11. <https://doi.org/10.1037/xge0000039>
- Wetzels, R., Vandekerckhove, J., Tuerlinckx, F., & Wagenmakers, E.-J. (2010). Bayesian parameter estimation in the expectancy valence model of the Iowa gambling task. *Journal of Mathematical Psychology*, *54*, 14–27. <https://doi.org/10.1016/j.jmp.2008.12.001>
- White, C. N., Servant, M., & Logan, G. D. (2017). Testing the validity of conflict drift-diffusion models for use in estimating cognitive

- processes: a parameter-recovery study. *Psychonomic Bulletin & Review*. <https://doi.org/10.3758/s13423-017-1271-2>.
- Worthy, D. A., Pang, B., & Byrne, K. A. (2013). Decomposing the roles of perseveration and expected value representation in models of the Iowa gambling task. *Frontiers in Psychology*, *4*, 1–9. <https://doi.org/10.3389/fpsyg.2013.00640>
- Yechiam, E., & Busemeyer, J. R. (2005). Comparison of basic assumptions embedded in learning models for experience-based decision making. *Psychonomic Bulletin & Review*, *12*, 387–402. <https://doi.org/10.3758/BF03193783>
- Yechiam, E., & Ert, E. (2007). Evaluating the reliance on past choices in adaptive learning models. *Journal of Mathematical Psychology*, *51*, 75–84. <https://doi.org/10.1016/j.jmp.2006.11.002>
- Yechiam, E., Stout, J. C., Busemeyer, J. R., Rock, S. L., & Finn, P. R. (2005). Individual differences in the response to forgone payoffs: An examination of high functioning drug abusers. *Journal of Behavioral Decision Making*, *18*, 97–110. <https://doi.org/10.1002/bdm.487>