



HAL
open science

The Principle of Redundant Reflection

Martin Metodiev, Maarten Marsman, Lourens Waldorp, Quentin Gronau, Eric-Jan Wagenmakers

► **To cite this version:**

Martin Metodiev, Maarten Marsman, Lourens Waldorp, Quentin Gronau, Eric-Jan Wagenmakers. The Principle of Redundant Reflection. 2025. <hal-05010871v2>

HAL Id: hal-05010871

<https://hal.science/hal-05010871v2>

Preprint submitted on 18 Sep 2025



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

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



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

The Principle of Redundant Reflection

Martin Metodiev ¹, Maarten Marsman ², Lourens Waldorp ²,
Quentin F. Gronau ³, and Eric-Jan Wagenmakers ²

Correspondence concerning this manuscript should be addressed to:

Martin Metodiev
Université Clermont Auvergne,
Laboratoire de Mathématiques Blaise Pascal
Campus des Cézeaux
3, place Vasarely
63178 Aubière cedex, France

E-mail may be sent to martin.metodiev@doctorant.uca.fr.

Abstract

The fact that redundant information does not update a rational belief implies that rational beliefs are updated using Bayes' rule. In the framework of Hild (1998a), this is true under mild conditions for discrete, continuous, and arbitrary measure spaces. We prove this result and illustrate it with two examples.

There exist multiple arguments as to why rational beliefs must be updated using the laws of probability (e.g., the Cox axioms, Cox, 1946; Dutch book arguments, Ramsey, 1926; de Finetti, 1931/1989; closeness-to-the-truth arguments, Joyce, 1998; see also Jeffreys, 1939; Carnap, 1950). Our goal is to add another such argument based on a simple and intuitive rule that we term *the principle of redundant reflection*.

The rule holds that once multiple accounts θ of the world have been completely specified and assigned prior probabilities $p(\theta)$, the mere contemplation of potential future outcomes Y^* (i.e., consideration of the prior predictive distribution) does not add novel information and therefore ought to leave the prior probabilities unaffected. In other words: the expected posterior probability (i.e., $\mathbb{E}_{Y^* \sim p(Y^*)} [p(\theta | Y^*)]$) equals the prior probability $p(\theta)$ (e.g., Chamley, 2004; Goldstein, 1983; Huttegger, 2017; Skyrms, 1997; Van Fraassen, 1984; see also Cook, Gelman, & Rubin, 2006; Fong, Holmes, & Walker, 2024; Gandy & Scott, 2021; Geweke, 2004; Schad, Nicenboim, Bürkner, Betancourt, & Vasishth, 2023; Talts, Betancourt, Simpson, Vehtari, & Gelman, 2018; Wagenmakers & Grasman, 2025; this is also known as *the iteration principle*, Hild, 1998a, 1998b; Spohn, 1978, p. 161;

¹Université Clermont Auvergne, Laboratoire de Mathématiques Blaise Pascal

²University of Amsterdam, Department of Psychology, Nieuwe Achtergracht 129B, 1018VZ Amsterdam, The Netherlands

³University of Newcastle, Australia

Spohn, 2012, p. 191). We refer to beliefs as rational if they are updated via this principle. Agents are rational if they use such an updating rule.

The principle of redundant reflection follows immediately from the law of total probability. Let each potential future realization of discrete data be denoted y_i^* . We then have

$$\begin{aligned}\mathbb{E}_{Y^* \sim p(Y^*)} [p(\theta | Y^*)] &= \sum_i p(\theta | y_i^*) p(y_i^*) \\ &= p(\theta).\end{aligned}$$

This result also holds for continuous probability density functions (replacing sums by integrals). The result is intuitive because the prior predictive Y^* is a direct consequence of the model definition: by specifying the model, the analyst indirectly also specifies the prior predictive distribution. If contemplating the prior predictive distribution would provide cause to adjust the prior probabilities, this can only mean that the original model specification was faulty or incomplete; a rational agent would have adjusted the models *ex ante*. In other words, the prior predictive is epistemically redundant given the model specification. In the statistical literature, the process of an unfolding posterior probability is known as a *martingale* and typifies a fair game (de Finetti, 1974, pp. 345–346; see also DeGroot, 1970; Doob, 1953, 1971; Lévy, 1937; Skyrms, 1997; Ville, 1939). As stated by Doob (1971, p. 454), “successive conditional expectations (...), as we know more and more, yield a martingale. (...) the game man plays with nature as he learns more and more is fair.”

The examples below highlight the delicate balance of factors that is required such that the expected posterior distribution equals the prior distribution. This balance is achieved automatically by applying the rules of probability.

Example 1: Contemplating an update to a posterior distribution

An analyst assigns a beta(2,2) prior distribution to a binomial chance parameter θ . The analyst contemplates the expected posterior distribution based on collecting $n^* = 2$ future observations. There are $n^* + 1 = 3$ possible number of successes s^* that may be observed. Under a beta(2,2) prior, the probability of outcomes $s^* = \{0, 1, 2\}$ equals $\{3/10, 4/10, 3/10\}$, respectively. Hence, the expected posterior distribution is a mixture of three beta distributions, with the mixture weights set equal to the respective outcome probabilities:

$$\begin{aligned}\mathbb{E}_{s^* \sim p(s^*)} [p(\theta | n^* = 2)] &= \frac{3}{10} \cdot p(\theta | s^* = 2) + \frac{4}{10} \cdot p(\theta | s^* = 1) + \frac{3}{10} \cdot p(\theta | s^* = 0) \\ &= \frac{3}{10} \cdot \text{beta}(\theta | 4, 2) + \frac{4}{10} \cdot \text{beta}(\theta | 3, 3) + \frac{3}{10} \cdot \text{beta}(\theta | 2, 4).\end{aligned}$$

This three-component mixture is equal to the beta(2,2) prior distribution. Had the analyst instead contemplated $n^* = 100$ future observations, the expected posterior would have been a 101-component mixture of beta distributions, whose weighted average would again equal the beta(2,2) prior distribution.

Example 2: Contemplating an update to a posterior model probability

An analyst entertains two models for a binomial chance parameter θ , the null hypothesis $\mathcal{H}_0: \theta = 1/2$ and the alternative hypothesis $\mathcal{H}_1: \theta \sim \text{beta}(1, 1)$. The analyst deems

\mathcal{H}_0 and \mathcal{H}_1 equally plausible *a priori*. The first five observations are all successes, and this gradually increases the probability for \mathcal{H}_1 from $1/2$ to $16/19 \approx 0.84$. The analyst anticipates collecting the sixth observation and wishes to compute the expected posterior probability for \mathcal{H}_1 .

Denote the hypothetical success and failure on the sixth trial by s^* and f^* , respectively. Note that the posterior odds for \mathcal{H}_1 over \mathcal{H}_0 equals $16/3 \approx 5.33$, and that the posterior distribution for θ under \mathcal{H}_1 is a beta(6,1) distribution. The probability that the sixth observation will be a success can be obtained by applying the law of total probability:

$$\begin{aligned} p(s^* | y) &= p(s^* | \mathcal{H}_0, y) \cdot p(\mathcal{H}_0 | y) + p(s^* | \mathcal{H}_1, y) \cdot p(\mathcal{H}_1 | y) \\ &= \frac{1}{2} \cdot \frac{3}{19} + \frac{6}{7} \cdot \frac{16}{19} \\ &= \frac{213}{266} \approx 0.80, \end{aligned}$$

and hence $p(f^* | y) = 53/266 \approx 0.20$. In case the sixth observation is a success, the associated Bayes factor BF_{10} is $(6/7)/(1/2) = 12/7$. The total posterior odds for \mathcal{H}_1 after this observation then equals $16/3 \cdot 12/7 = 192/21$, which translates to a posterior probability of $(192/21)/(192/21 + 1) = 192/213$. In case the sixth observation is a failure, the associated Bayes factor BF_{10} is $(1/7)/(1/2) = 2/7$. The total posterior odds for \mathcal{H}_1 after this observation then equals $16/3 \cdot 2/7 = 32/21$, which translates to a posterior probability of $(32/21)/(32/21 + 1) = 32/53$.

In sum: (a) there is a probability of $213/266 \approx 0.80$ of observing a success, which then yields a posterior probability for \mathcal{H}_1 of $192/213 \approx 0.90$ (i.e., an increase of about 0.06); (b) there is a probability of $53/266 \approx 0.20$ of observing a failure, which then yields a posterior probability for \mathcal{H}_1 of $32/53 \approx 0.60$ (i.e., a decrease of about 0.24). When contemplating the updated prior predictive $p(s^* | y)$, the expected posterior probability for \mathcal{H}_1 therefore equals

$$\begin{aligned} \mathbb{E}_{s^* \sim p(s^* | y), f^* \sim p(f^* | y)} [p(\mathcal{H}_1 | s^*, f^*, y)] &= p(\mathcal{H}_1 | s^*, y) \cdot p(s^* | y) + p(\mathcal{H}_1 | f^*, y) \cdot p(f^* | y) \\ &= \frac{192}{213} \cdot \frac{213}{266} + \frac{32}{53} \cdot \frac{53}{266} \\ &= \frac{16}{19} = p(\mathcal{H}_1 | y). \end{aligned}$$

Observing a sequence of five successes gradually strengthens the expectation that the next observation will also be a success, in which case the probability for \mathcal{H}_1 would continue to increase; however, there remains a probability that the next trial will be a failure, and this will greatly decrease the probability for \mathcal{H}_1 . The probabilities for the possible outcomes and the effect that these outcomes will have on the posterior plausibilities of the models *cancel out exactly* so that the expected model probability equals the prior model probability.

To conclude: when contemplating hypothetical future outcomes, the analyst generally expects to gain information, but is completely uncertain about the impact that this future information will have on the model probability of interest. This holds irrespective of how deeply the analyst looks into the future (e.g., the result holds regardless of whether the expected posterior probability is computed across the next observation, the next 100 observations, or the next 100 billion observations). The result also holds at any point in time as the observations accumulate (e.g., the analyst continually updates their prior to a posterior; because the posterior acts as a prior for the next batch of data, the informational impact of the new batch is completely unknown).

The Necessity of Bayes' Rule

The fact that redundant information should not update a rational belief follows from the law of total probability. We now show that the law of total probability implies the involvement of Bayes' rule in updating beliefs. First note that the principle of redundant reflection is consistent with an update of beliefs that adheres to Bayes' rule:

$$\begin{aligned}
 \mathbb{E}_{Y^* \sim p(Y^*)} [p(\theta | Y^*)] &= \sum_i p(\theta | y_i^*) p(y_i^*) \\
 &= \sum_i \left[p(\theta) \frac{p(y_i^* | \theta)}{p(y_i^*)} \right] p(y_i^*) \\
 &= p(\theta) \sum_i \frac{p(y_i^* | \theta)}{p(y_i^*)} p(y_i^*) \\
 &= p(\theta) \sum_i p(y_i^* | \theta) \\
 &= p(\theta).
 \end{aligned} \tag{1}$$

The first step in this equation expands the posterior $p(\theta | y_i^*)$ using Bayes' rule. Next, the terms $p(y_i^*)$ cancel, the term $p(\theta)$ (i.e., the desired answer) is placed in front of the summation, and the remaining summation runs over all possible data y_i^* and therefore equals 1.

We now wish to show that using Bayes' rule to expand the posterior $p(\theta | y_i^*)$ constitutes the only reasonable way that the law of total probability will hold (and, consequently, the only way that the expected posterior probability equals the prior probability). In other words, we wish to show that the principle of redundant reflection is *uniquely* consistent with an update of beliefs that adheres to Bayes' rule.

We will employ the framework of Hild (1998a) for our proof. This framework is well grounded in the literature (i.e., Hild, 1998a, 1998b; Spohn, 2012, Chapter 9); correspondence to the conditions given in Hild (1998a) will be indicated explicitly. We believe these conditions are straightforward and innocuous, but a discussion on this point is delayed until after the full statement.

Let P be a probability measure defining the joint distribution of Y^*, ϑ over $\mathcal{Y} \times \Theta$, where Y^* is a random variable taking values in a measurable space (\mathcal{Y}, Σ_Y) , ϑ is a random variable taking values in a measurable space $(\Theta, \Sigma_\vartheta)$. Furthermore, let Y^* denote the data and y^* its realization, and let ϑ denote the parameter and θ its realization. In addition, we use $p(\cdot | B)$, and $p(\cdot)$ to denote the conditional and marginal Radon-Nikodym derivatives of Y^* with respect to the dominating measure λ (note that in practice these are usually either probability mass functions or probability density functions). It is assumed that these derivatives exist and are well-defined for any non-zero event.

We assume that a prior belief in a state of the world is quantified by a single probability measure P . This means that any event that one can have beliefs about is characterized by a set A , and $P(A)$ is the prior belief in A . Note that we are using the notion of prior belief in its broadest sense: one has a prior belief in the parameter, but this belief in the parameter also defines one's prior belief in the data through the prior predictive. We use this specific

measure-theoretic notation to reflect this fact, since A can be any event that one has beliefs about, whether A involves the data or the parameter. Our prior belief function $\text{PriorBel}(A)$ is thus given by

$$\text{PriorBel}(A) = P(A) \quad \text{for any } P\text{-measurable set } A. \quad (2)$$

This corresponds to the *auto-epistemic transparency condition* in Hild (1998a). The updated probability measure UpdatedBel_{Y^*} is assigned based on the data Y^* , and it is assumed that the change in belief brought about by the data is entirely determined by the data itself, in the sense that

$$\text{UpdatedBel}_{Y^*}(A) \text{ is } Y^*\text{-measurable for any } P\text{-measurable set } A. \quad (3)$$

This implies, for instance, that other random variables that have nothing to do with the data do not change the updated probability measure, since this probability measure was updated based only on the data, and nothing else. This is identical to the rule of *evidence-driven updating* described in Hild (1998a).

There is no restriction on UpdatedBel_{Y^*} (it does, e.g., not have to follow Bayes' rule), other than it taking the data Y^* as certain, meaning that for any realization y^* of Y^*

$$\text{UpdatedBel}_{y^*}(A) = \begin{cases} 1 & \{Y^* = y^*\} \subseteq A, \\ 0 & \{Y^* = y^*\} \subseteq \bar{A}, \end{cases} \quad \text{for any } P\text{-measurable set } A, \quad (4)$$

with \bar{A} denoting the complement of A . It should thus be noted that we take the realization of the data y^* as exactly what it is, without conveying any additional information. In particular, we are completely certain of the information $\{Y^* = y^*\}$, and the information $\{Y^* = y^*\}$ cannot be self-contradictory in the sense that it will never lead us to believe that $\{Y^* = y^*\}$ is wrong. This corresponds to the *crux condition* in Hild (1998a).

Under these conditions, we define the principle of redundant reflection as follows:

Definition 1 (the principle of redundant reflection) *The updated belief UpdatedBel_{Y^*} follows the principle of redundant reflection if and only if it satisfies*

$$\text{PriorBel}(A) = E_{Y^* \sim p(Y^*)}[\text{UpdatedBel}_{Y^*}(A)] \quad \text{for any } P\text{-measurable set } A. \quad (5)$$

Perhaps surprisingly, the principle of redundant reflection implies that beliefs are updated using Bayes' rule.

Theorem 1 *Suppose that the prior belief is defined by auto-epistemic transparency (Equation 2). Then, any updated belief UpdatedBel_{Y^*} that is evidence driven (Equation 3), meets the crux condition (Equation 4), and obeys the principle of redundant reflection (Definition 1) is defined via Bayes' rule:*

$$\text{UpdatedBel}_{Y^*}(B) = \frac{\text{PriorBel}(B) p(Y^* | B)}{p(Y^*)} \quad P\text{-almost surely} \quad (6)$$

for any P -measurable B with non-zero probability. In particular, Equation (6) holds if $B = \{\vartheta \in D\}$, where D is a ϑ -measurable event for which $P(\vartheta \in D) > 0$.

The proof is given in the appendix.

Concluding Comments

The claim that redundant information should leave a rational belief unaffected appears intuitive and innocuous, and follows directly from the law of total probability. Under mild conditions, however, this law of total probability implicitly assumes that beliefs are updated using Bayes' rule. In other words, once an agent assigns a prior distribution across rival accounts of the world θ , this agent is then forced by the principle of redundant reflection to update that probability distribution exclusively using Bayes' rule.

To the best of our knowledge, the only case where a similar argument has been made before is Hild (1998a). However, the argument was shown only for discrete distributions and was made in a somewhat roundabout way, showing that the principle of redundant reflection (named *the principle of iteration* in Hild, 1998a) is equivalent to the principle of reflection (Van Fraassen, 1984, p. 244), the principle of reflection is equivalent to auto-epistemic updating for evidence-driven updating rules, and auto-epistemic updating is equivalent to conditioning (i.e., the application of Bayes' rule) if the crux condition is fulfilled. We aimed to clarify the issue by showing the equivalence between the principle of redundant reflection and Bayes' rule directly. We also demonstrated equivalence in arbitrary measure spaces. In particular, this is useful when the data are continuous. While Hild mentions that this extension of his result is likely possible, the proof from Hild (1998a) does not automatically apply here, since all point-events have probability 0.

It should also be noted that, while our result is an extension of the result of Hild (1998a) in the mathematical sense, it is different in the philosophical sense. The auto-epistemic paradigm focuses on what one must believe now, given that one knows what one will believe in the future. We use the same probabilistic phenomenon, but in a slightly different way: we state that given the complete model specification, there is no additional information in the predictions, so examining the predictions should leave one's beliefs fully intact. So one does not need to assess what one's beliefs will be in the future – the beliefs are assessed now, and one realizes that the predictions are redundant and do not convey new information. So, the current belief state constrains the set of future belief states, whereas in Hild (1998a) it is presented more the other way around: future belief states constrain the current belief state.

There are two caveats to our uniqueness result: the first is that it only shows uniqueness of Bayes' rule *almost surely*. There exist updating rules that satisfy all our assumptions, but differ from Bayes' rule on a set with probability 0. The second caveat is that our result makes use of the assumptions that the prior belief is characterized by a probability measure, that any updating rule must be evidence driven and that the evidence must be taken as absolutely certain, delivering no other information than the evidence itself. While we consider the former assumptions innocuous, the latter is somewhat controversial in an epistemological sense (see Vineberg, 2022, Section 4.3 for a review). However, in a scientific setting this assumption is a lot more reasonable, since the Bayesian model can always be expanded to include any uncertainty about the data. In fact, as Hild (1998b) points out, the framework that the principle of reflection is based on is a framework “which consciously confines itself to ideally rational agents” (Hild, 1998b, page 2). In our opinion, this is exactly the framework that science should operate in.

Acknowledgements

The authors would like to thank Nicholas J. Irons for useful discussions on posterior distributions, as well as Alexander Reisach and Niklas Jacobs for helpful comments.

References

- Carnap, R. (1950). *Logical foundations of probability*. Chicago: The University of Chicago Press.
- Chamley, C. P. (2004). *Rational herds: Economic models of social learning*. Cambridge University Press: Cambridge.
- Cook, S. R., Gelman, A., & Rubin, D. B. (2006). Validation of software for Bayesian models using posterior quantiles. *Journal of Computational and Graphical Statistics*, 15, 675–692.
- Cox, R. T. (1946). Probability, frequency and reasonable expectation. *The American Journal of Physics*, 14, 1–13.
- de Finetti, B. (1931/1989). Probabilism: A critical essay on the theory of probability and on the value of science. *Erkenntnis*, 31, 169–223.
- de Finetti, B. (1974). *Theory of probability, vol. 1 and 2*. New York: John Wiley & Sons.
- DeGroot, M. H. (1970). *Optimal statistical decisions*. New York: McGraw–Hill.
- Doob, J. L. (1953). *Stochastic processes*. New York: John Wiley & Sons.
- Doob, J. L. (1971). What is a martingale? *The American Mathematical Monthly*, 78, 451–463.
- Fong, E., Holmes, C. C., & Walker, S. G. (2024). Martingale posterior distributions. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 85, 1357–1391.
- Gandy, A., & Scott, J. A. (2021). Unit testing for MCMC and other Monte Carlo methods. *ArXiv Preprint*. Retrieved from <https://arxiv.org/abs/2001.06465>
- Geweke, J. (2004). Getting it right: Joint distribution tests of posterior simulators. *Journal of the American Statistical Association*, 99, 799–804.
- Goldstein, M. (1983). The prevision of a prevision. *Journal of the American Statistical Association*, 78, 817–819.
- Hild, M. (1998a). Auto–epistemology and updating. *Philosophical Studies*, 92, 321–361.
- Hild, M. (1998b). The coherence argument against conditionalization. *Synthese*, 115, 229–258.
- Howson, C., & Urbach, P. (2006). *Scientific reasoning: The Bayesian approach (3rd. ed.)*. Chicago: Open Court.
- Huttegger, S. M. (2017). *The probabilistic foundations of rational learning*. Cambridge: Cambridge University Press.
- Jeffreys, H. (1939). *Theory of probability* (1st ed.). Oxford, UK: Oxford University Press.
- Joyce, J. M. (1998). A non–pragmatic vindication of probabilism. *Philosophy of Science*, 65, 575–603.
- Lévy, P. (1937). *Théorie de l'addition des variables aléatoires*. Paris: Gauthier-Villars.
- Ramsey, F. P. (1926). Truth and probability. In R. B. Braithwaite (Ed.), *The foundations of mathematics and other logical essays* (pp. 156–198). London: Kegan Paul.
- Schad, D. J., Nicenboim, B., Bürkner, P.-C., Betancourt, M., & Vasishth, S. (2023). Workflow techniques for the robust use of Bayes factors. *Psychological Methods*, 28, 1404–1426.
- Skyrms, B. (1997). The structure of radical probabilism. *Erkenntnis*, 45, 285–297.
- Spohn, W. (1978). *Grundlagen der Entscheidungstheorie*. Kronberg/Ts.: Scriptor–Verlag.
- Spohn, W. (2012). *The laws of belief: Ranking theory and its philosophical applications*. Oxford: Oxford University Press.
- Talts, S., Betancourt, M., Simpson, D., Vehtari, A., & Gelman, A. (2018). Validating Bayesian inference algorithms with simulation–based calibration. *ArXiv Preprint*. Retrieved from <https://arxiv.org/abs/1804.06788>
- Van Fraassen, B. C. (1984). Belief and the will. *The Journal of Philosophy*, 81, 235–256.
- Ville, J. (1939). *Étude critique de la notion de collectif* (Unpublished doctoral dissertation). La Faculté des Sciences de Paris.
- Vineberg, S. (2022). Dutch book arguments. In E. N. Zalta & U. Nodelman (Eds.), *The stanford encyclopedia of philosophy*. Metaphysics Research Lab, Stanford University. Retrieved from <https://plato.stanford.edu/archives/fall12022/entries/dutch-book/>
- Wagenmakers, E.-J., & Grasman, R. P. P. P. (2025). A discrepancy measure based on expected

posterior probability. *PsyArXiv*. Retrieved from https://osf.io/preprints/psyarxiv/zjmsq_v1

Appendix: Proof of Theorem 1

Proof. For any $C \in \Sigma_Y$ for which $P(Y^* \in C)$ is non-zero the principle of redundant reflection, inserting $A = \{B, Y^* \in C\}$, implies that

$$\text{PriorBel}(B, Y^* \in C) = \int \text{UpdatedBel}_{y^*}(B, Y^* \in C)p(y^*) d\lambda(y^*),$$

where the integral is to be understood measure-theoretically, being replaced by a sum if Y^* is discrete.

The crux condition, inserting $A = \{B, Y^* \in C\}$, implies that

$$\text{UpdatedBel}_{y^*}(B, Y^* \in C) = 0 \quad \text{if} \quad y^* \notin C,$$

since $\{Y^* = y^*\} \subseteq \overline{\{Y^* \in C\}}$ in that case. Thus

$$\text{PriorBel}(B, Y^* \in C) = \int_C \text{UpdatedBel}_{y^*}(B, Y^* \in C)p(y^*) d\lambda(y^*). \quad (7)$$

Now note that $\text{UpdatedBel}_{y^*}(A)$ is defined for any event A involving Y^* or ϑ . It follows that A can be equal to the marginal event $\{Y^* \in C\}$, since $\{Y^* \in C\} = \{\vartheta \in \Theta, Y^* \in C\}$ is also an event involving ϑ or Y^* , although there is no restriction on ϑ . In addition, $y^* \in C$ implies that $\{Y^* = y^*\} \subseteq \{Y^* \in C\}$, so by the crux condition, inserting $A = \{Y^* \in C\}$,

$$\text{UpdatedBel}_{y^*}(Y^* \in C) = 1.$$

Since UpdatedBel_{y^*} is a probability measure it follows that events with probability 1 do not change this measure. Hence

$$\text{PriorBel}_{y^*}(B, Y^* \in C) = \int_C \text{UpdatedBel}_{y^*}(B)p(y^*) d\lambda(y^*).$$

Notice that if Bel_{y^*} was a conditional probability measure, this would also be true, since we would be able to use result (8) in Howson and Urbach (2006).

Dividing and multiplying the right-hand side by $P(Y^* \in C)$ gives

$$\text{PriorBel}_{y^*}(B, Y^* \in C) = \frac{\int_C \text{UpdatedBel}_{y^*}(B)p(y^*) d\lambda(y^*)}{P(Y^* \in C)} P(Y^* \in C).$$

The expression next to $P(Y^* \in C)$ is the same as the expectation of $\text{UpdatedBel}_{Y^*}(B)$, conditional on $Y^* \in C$. It follows that Equation (7) is equivalent to

$$\text{PriorBel}(B, Y^* \in C) = E_{Y^* \sim P(Y^*|Y^* \in C)}[\text{UpdatedBel}_{Y^*}(B)]P(Y^* \in C). \quad (8)$$

Dividing both ends by $P(Y^* \in C)$ and replacing PriorBel by P (Equation 2) gives

$$\frac{P(B, Y^* \in C)}{P(Y^* \in C)} = E_{Y^* \sim P(Y^*|Y^* \in C)}[\text{UpdatedBel}_{Y^*}(B)]. \quad (9)$$

In other words, the ratio of the prior probability of the joint events $\{B, Y^* \in C\}$ and the prior probability of the marginal event $\{Y^* \in C\}$ is equal to the conditional expectation.

Note that Equation (9) was thus shown for an arbitrary measurable set C whose probability is not 0. It is also easy to see that the principle of redundant reflection is indeed fulfilled by the probability measure defined in Equation (6). Thus, we only need to show that there is no other probability measure $\text{UpdatedBel}'$ fulfilling Equation (6) and all of the other conditions outlined in Theorem 1. We can do this by showing that

$$E_{Y^* \sim p(Y^*)}[|\text{UpdatedBel}'_{Y^*}(B) - \text{UpdatedBel}_{Y^*}(B)|] = 0,$$

since equality in expected absolute difference implies almost sure equality.

Indeed, we can use Equation (9), once by inserting the set $C = \zeta = \{Y^* : \text{UpdatedBel}'_{Y^*}(B) > \text{UpdatedBel}_{Y^*}(B)\}$ and again by inserting $C = \bar{\zeta}$ (measurable since both updating rules are evidence driven by assumption), and the law of total probability to obtain

$$\begin{aligned} & E_{Y^* \sim p(Y^*)}[|\text{UpdatedBel}'_{Y^*}(B) - \text{UpdatedBel}_{Y^*}(B)|] = \\ & P(Y^* \in \zeta) E_{Y^* \sim p(Y^*|Y^* \in \zeta)}[\text{UpdatedBel}'_{Y^*}(B) - \text{UpdatedBel}_{Y^*}(B)] + \\ & P(Y^* \in \bar{\zeta}) E_{Y^* \sim p(Y^*|Y^* \in \bar{\zeta})}[\text{UpdatedBel}'_{Y^*}(B) - \text{UpdatedBel}_{Y^*}(B)] = \\ & P(Y^* \in \zeta)(E_{Y^* \sim p(Y^*|Y^* \in \zeta)}[\text{UpdatedBel}'_{Y^*}(B)] - E_{Y^* \sim p(Y^*|Y^* \in \zeta)}[\text{UpdatedBel}_{Y^*}(B)]) + \\ & P(Y^* \in \bar{\zeta})(E_{Y^* \sim p(Y^*|Y^* \in \bar{\zeta})}[\text{UpdatedBel}'_{Y^*}(B)] - E_{Y^* \sim p(Y^*|Y^* \in \bar{\zeta})}[\text{UpdatedBel}_{Y^*}(B)]) = 0, \end{aligned}$$

in the case that $P(Y^* \in \zeta) \in (0, 1)$. The cases $P(Y^* \in \zeta) = 0$ and $P(Y^* \in \zeta) = 1$ follow analogously. This completes the proof. \square