Volume 35, Issue 1

Expanding the Weighted Updating Model

Jesse Aaron Zinn

College of Business, Clayton State University

Abstract

This work casts light upon a pair of restrictions inherent to the basic weighted updating model, which is a generalization of Bayesian updating that allows for biased learning. Relaxing the restrictions allows for the study of individuals who discriminate between observations or who treat information in a dynamically inconsistent manner. These generalizations augment the set of cognitive biases that can be studied using new versions of the weighted updating model to include the availability heuristic, order effects, self-attribution bias, and base-rate neglect in light of irrelevant information.

I am grateful for valuable comments and suggestions from Ted Bergstrom, Javier Birchenall, Gary Charness, Zack Grossman, Jason Lepore, Dick Startz, an anonymous referee, and seminar discussants at Cal Poly San Luis Obispo, the 2013 conference of the Society for the Advancement of Behavior Economics, and the 2014 Bay Area Behavioral and Experimental Economics Workshop. Any errors are to be attributed to the author.

Citation: Jesse Aaron Zinn, (2015) "Expanding the Weighted Updating Model", *Economics Bulletin*, Volume 35, Issue 1, pages 182-186

Contact: Jesse Aaron Zinn - JesseZinn@clayton.edu.

Submitted: September 14, 2014. **Published:** March 11, 2015.

1. Introduction

The weighted updating model generalizes Bayesian updating to allow for irrational learning by exponentially weighting the likelihood function $f(h_t|\theta)$ and the prior $\pi(\theta)$ as in

$$\tilde{\pi}(\theta|h_t) = \frac{f(h_t|\theta)^\beta \pi(\theta)^\alpha}{\int_{\Theta} f(h_t|\theta)^\beta \pi(\theta)^\alpha d\theta}, \quad (1)$$

where h_t denotes an ordered history of observations $(x_1, \dots, x_t)$ that are used to infer the value of the distribution parameter $\theta \in \Theta$. It is straightforward to see that Bayes' rule is the special case of (1) with $\alpha = \beta = 1$. Intuitive interpretations of the weights are relatively straightforward: When $\beta < 1$ the history of observations h_t is effectively being treated as less informative than a perfect Bayesian would treat h_t (and vice versa). An analogous interpretation holds for α with respect to prior information.¹ For example, conservatism bias describes beliefs that do not change to high enough degree in light of new evidence. This bias can be modelled using (1) by $\alpha \geq 1$ and $\beta \leq 1$, as long as at least one of the inequalities is strict.

The purpose of this work is to point out and show how to overcome two restrictions inherent to the basic weighted updating model expressed in (1). The first restriction is that the same weight β measures the information content of all of the observed outcomes $x_1, \dots, x_t$. In Section 2, I use the definition of conditional distribution functions to motivate a version of (1) where each x_j has its own weight β_j . Modelling beliefs in this way allows for the study of individuals who discriminate between observations. The second restriction implicit to the basic weighted updating model (1) is that the weights are fixed over time, as additional data are observed. In Section 3, I explore the possibility of allowing α and β to vary over time. Relaxing either of these restrictions expands the set of biases to which the weighted updating model can be applied, adding at least the availability heuristic, order effects, self-attribution bias, and base-rate neglect in light of irrelevant information.

This paper builds upon the literature that utilizes the basic weighted updating model in expression (1). This literature includes Grether (1980) and Grether (1992), both of which provide empirical evidence for the representativeness heuristic by estimating the weights on the likelihood function and the prior distribution. Palfrey and Wang (2012) use weighted updating to model investors who under- or overreact to public information regarding financial assets in a model with speculative pricing. Benjamin *et al.* (2013) model “non-belief in the law of large numbers” (both on its own and in combination with the “law of small numbers” and base-rate neglect) using the weighted updating model.²

¹Zinn (2015) shows that the weights α and β systematically affect the information entropy of the distribution they are weighting, with greater weight yielding less information entropy. This provides a rigorous foundation for the intuitive interpretations mentioned above.

²Other related work includes Ibrahim and Chen (2000), which introduces *power priors*, a framework that allows the statistician to consider data from previous studies by finding a weight in $(0, 1)$ to put on that data while maintaining a weight of 1 on current data. This can be viewed as a case of weighted updating wherein the statistician rationally weights different batches of data. Van Benthem *et al.* (2009) define a “weighted product updating rule” and show that Bayes' rule and the Jeffrey updating rule are both special cases.

2. Discrimination Between Data

By the definition of conditional distribution functions, for any $t \in \mathbb{N}$ and likelihood function $f(h_t|\theta)$ the following decomposition holds:

$$f(h_t|\theta) = f(x_t|h_{t-1}, \theta)f(h_{t-1}|\theta). \quad (2)$$

Here it may be useful to note that if the observations are (or at least are assumed to be) independent then the history h_{t-1} pre-dating observation x_t does not affect its likelihood function, in which case $f(x_t|\theta)$ can replace $f(x_t|h_{t-1}, \theta)$ in expression (2) and throughout the paper.

Repeated iteration of expression (2) yields

$$f(h_t|\theta) = \prod_{j=1}^t f(x_j|h_{j-1}, \theta),$$

which motivates setting up the weighted updating model as³

$$\tilde{\pi}(\theta|h_t) \propto \pi(\theta)^\alpha \prod_{j=1}^t f(x_j|h_{j-1}, \theta)^{\beta_j}, \quad (3)$$

where α remains the weight on the prior distribution and β_j is the weight associated with the j th datum x_j , for each $j \in \{1, \dots, t\}$. Of course, the introductory framework in expression (1') is the special case of (3) wherein $\beta_j = \beta$ for each $j \in \{1, \dots, t\}$.

In light of the information-theoretic interpretations provided in Zinn (2015), we can say that if an individual's beliefs evolve according to the weighted updating model in (3), then, compared to a perfect Bayesian, the individual is subjectively treating the component distributions proportional to $\pi(\theta)^\alpha$ and $f(x_j|h_{j-1}, \theta)^{\beta_j}$ for $j = 1, \dots, t$ each as containing either more or less information content depending on how the levels of $\alpha, \beta_1, \dots, \beta_t$ compare to unity. As the prior $\pi(\theta)$ summarizes prior information and each likelihood function $f(x_j|h_{j-1}, \theta)$ represents the influence of an individual datum x_j , the weighted updating model in expression (3) essentially allows the individual to treat the prior information and each datum x_j at their own idiosyncratic levels of information content.

Additional biases that the new version of the weighted updating model expressed in (3) is capable of modelling include the availability heuristic; order effects, such as primacy and recency; and self-attribution bias. The remainder of this section discusses how to model these biases with weighted updating. Table 1 summarizes this discussion.

The availability heuristic generates biases due to certain observations being more available in memory (Tversky and Kahneman, 1973). This can be modelled using weighted

³Note that the marginal distribution in (the denominator of) expression (1) serves only to normalize the weighted posterior $\tilde{\pi}(\theta|h_t)$. Hence, the model can be expressed more simply by

$$\tilde{\pi}(\theta|h_t) \propto f(h_t|\theta)^\beta \pi(\theta)^\alpha. \quad (1')$$

Expression (3) makes use of this economy of notation.

Table 1: Biases Involving Discrimination between Non-Prior Data

Cognitive Bias	Weights
Availability	β_j high for x_j that are more salient
Primacy Effect	β_j decreasing in j
Recency Effect	β_j increasing in j
Self-Attribution	β_j low if x_j is undesirable

updating simply by assuming that an economic agent has greater values of β_j corresponding to observations x_j that are relatively memorable.

Order effects occur when the relative temporal position of observations seems to affect beliefs formed from those observations. Experimental subjects typically exhibit either the primacy effect, where earlier observations are more salient than later observations, or the recency effect, where the opposite occurs (Hogarth and Einhorn, 1992). To model the primacy effect with the weighted updating model would require that β_j decreases as j rises, while modelling the recency effect involves assuming that β_j is increasing in j .

Self-attribution bias occurs when individuals credit their own ability for desirable outcomes but blame undesirable outcomes on external factors, such as luck.⁴ This suggests that agents put greater weights on x_j that are desirable and lower weights on x_j that are undesirable.

3. Dynamically Inconsistent Weights

This paper has, up to this point, presented the weighted updating model as one in which the weights are fixed. This section discusses relaxing this restriction so that weights can change over time.

To allow the weights to change over time involves allowing them to be functions of time. These functions can be defined exogenously or endogenously depending on the nature of the application. Denote them with $\alpha(t)$ and $\beta(t)$, so that after observing h_t the basic weighted updating model in expression (1') becomes

$$\tilde{\pi}(\theta|h_t) \propto f(h_t|\theta)^{\beta(t)}\pi(\theta)^{\alpha(t)}.$$

A bias that can be modelled with weights that change over time is base-rate neglect in light of irrelevant information. As its name suggests, base-rate neglect involves ignoring prior information, which in its most extreme form can simply be modelled by setting $\alpha = 0$, as is mentioned in the appendix of Benjamin *et al.* (2013). However, subjects who exhibit

⁴Self-attribution bias seems to depend also on a multitude of factors, including the mood of the individual, their type of culture, and the social setting, all of which could be taken into account by adjusting the weights in the weighted updating model. See Shepperd *et al.* (2008) for a recent survey of this literature.

base-rate neglect in light of irrelevant information typically do not ignore prior information until after they have observed some non-prior information, suggesting $\alpha > 0$ before observing non-prior information. This is illustrated in an experiment on base-rate neglect described in Kahneman and Tversky (1973). In this experiment, base rates differed between subjects: One group was told that the descriptions they observed were drawn from a population of 70 lawyers and 30 engineers, while the other group was told that they were drawn from a population with the frequencies reversed, 30 lawyers and 70 engineers. When experimental subjects observed a purposefully uninformative description of a man and were asked to guess whether he is an engineer or a lawyer, the average guess at the probability that the man was an engineer was approximately 50% *in both groups*. This base-rate neglect occurred even though the likelihoods participants gave were consistent with base rates before observing the irrelevant information, suggesting that participants utilized base rates then ignored them after observing the uninformative description. Such a phenomenon can be modelled by defining $\alpha(t)$ such that $\alpha(0) > 0$ (so that agents utilize prior information) and $\alpha(t) = 0$ for $t > 0$ (so that they ignore the prior information after observing any history h_t).⁵

Bar-Hillel (1980) finds evidence that individuals who exhibit base-rate neglect do so because prior information loses salience once there are non-prior observations to consider. These findings arose from experiments in which subjects made inferences based solely on prior information several times before observing non-prior information. Bar-Hillel hypothesized that this made the prior information more salient to participants. Modelling this with base-rate neglect might involve assuming that α is a function that is decreasing in t (at least from $t = 0$ to $t = 1$), but increasing in the number of times prior information was used before observations are made (i.e. when $t = 0$).

4. Concluding Remarks

In this paper I have shown how to relax two implicit restrictions of the basic weighted updating model. Eliminating these restrictions augments the list of biases that weighted updating can model. That list now includes the availability heuristic, base-rate neglect, the law of small numbers, non-belief in the law of large numbers, order effects (e.g. recency and primacy), the representativeness heuristic, and self-attribution bias. The author does not expect that this list exhausts the set of biases that can be modelled with weighted updating.

References

- Bar-Hillel M. (1980) “The Base-Rate Fallacy in Probability Judgments” *Acta Psychologica* **44**, 211–233.
- Benjamin D.J., Rabin M., and Raymond C. (2013) “A Model of Non-Belief in the Law of Large Numbers” Oxford University Department of Economics Discussion Paper No. 672.

⁵Rabin (1996) points out that weighted updating with constant weights cannot account for base-rate neglect in light of irrelevant information. (See footnote 60 of Rabin (1996). Note that the version of this paper published in 1998 in *Journal of Economic Literature* does not include this discussion.)

- Grether D.M. (1980) “Bayes Rule as a Descriptive Model: The Representativeness Heuristic” *Quarterly Journal of Economics* **95**, 537–557.
- Grether D.M. (1992) “Testing Bayes’ Rule and the Representativeness Heuristic: Some Experimental Evidence” *Journal of Economic Behavior & Organization* **17**, 31–57.
- Hogarth R.M. and Einhorn H.J. (1992) “Order Effects in Belief Updating: The Belief-Adjustment Model” *Cognitive Psychology* **24**, 1–55.
- Ibrahim J.G. and Chen M.H. (2000) “Power Prior Distributions for Regression Models” *Statistical Science* **15**, 46–60.
- Kahneman D. and Tversky A. (1973) “On the Psychology of Prediction” *Psychological Review* **80**, 237–251.
- Palfrey T.R. and Wang S.W. (2012) “Speculative Overpricing in Asset Markets with Information Flows” *Econometrica* **80**, 1937–1976.
- Rabin M. (1996) “Psychology and Economics” UC Berkeley Working Paper.
- Shepperd J., Malone W., and Sweeny K. (2008) “Exploring Causes of the Self-serving Bias” *Social and Personality Psychology Compass* **2**, 895–908.
- Tversky A. and Kahneman D. (1973) “Availability: A Heuristic for Judging Frequency and Probability” *Cognitive Psychology* **5**, 207–232.
- Van Benthem J., Gerbrandy J., and Kooi B. (2009) “Dynamic Update with Probabilities” *Studia Logica* **93**, 67–96.
- Zinn J.A. (2015) “Modelling Biased Judgement with Weighted Updating” MPRA Paper No. 61403.