I Plenitude — a standard line

Confidences, or credences, come in different degrees which can be modeled by a probability function if they are rational.

Result 1 (Always update — when cost free)**.** We should, from the point of view of maximizing expected value, update on any information that we come across given (among other things) that *updating is cost free*.
(From I.J. Good's 1967 and Graham Oddie's 1997 results.)

II Scarcity — a realistic caveat

We have limited processing power and working memory that has to be divided up between cognitive endeavors. Updating incurs cognitive opportunity costs.

III Information response policies

The policies of interest are *ways* of responding to available information.

- It is not required that we be able to implement them — our responses to information tend to be automatic.
- We need not be aware of the policies we adhere to.
- They may or may not be policies that it would be good to *try* to adhere to — the effort involved in trying to adhere to them might be too cognitively and epistemically costly.

III.1 Two policies

Policy 1 (The naïve policy)**.** Take any piece of available information into account for further processing when sufficient cognitive resources are available.

Policy 2 (The obstinate policy)**.** Disregard any available information as bearing on a proposition once substantial information for that proposition has been processed, otherwise proceed naïvely.

IV Modeling features of our epistemic scenario

IV.1 Measuring accuracy

Letting 1 represent the truth value of a true proposition, 0 the truth value of a false proposition, and credences range over the unit interval $[0, 1]$, we judge a credence in p to being more accurate as the distance between that credence and the truth value of p decreases. For specificity, we will use a version of the popular Briar-score to measure inaccuracy. So, for a total credal state S with respect to each proposition p_i in $\{p_1, \ldots, p_n\}$:

$$D(S) = \sum_i (1 - \mathbb{P}_t[p_i])^2,$$

with $\mathbb{P}_t$ the subject's credence function at t.

IV.2 Veridicality and diminishing returns

I will assume that (i) our situation is the non-skeptical one and (ii) there are diminishing returns on learning. This will be captured by specifying how credences respond to information. Figure 2 displays a (modest) possible mapping that will be used in the case study.*

Figure 2. Normal updating of confidence in a truth.†

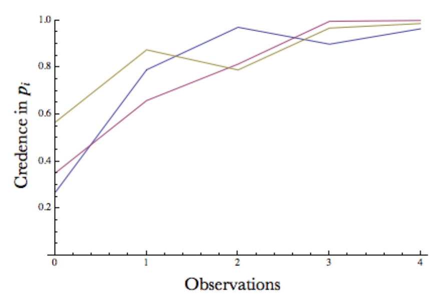

*The result is robust under functions that accord with these motivating premises.
†Here the credence function in the truth increases at a logistic rate with normally distributed variation.

IV.3 A queue-theoretic model of learning under processing and memory limitations

The mathematics of queues can help us understand how memory and processing constraints interact with revision policies to yield credences.

Figure 1. The general model of epistemic inquiry

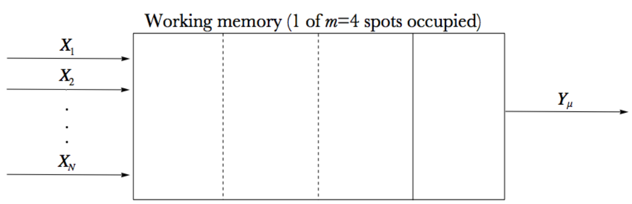

IV.4 The policies made formal

V Filling in the model

V.1 The distributions

In the simple case study, we will focus on situations in which (i) any relevant event is as likely to occur within any time interval of equal length, and (ii) whether or not an interval in question was recently preceded by another incident.[‡]

V.2 Manipulating the model's parameters

1. The amount of working memory m;
2. The number of questions of interest n; and
3. The rates of information arrival and processing.

V.3 The results

Result 2 (When obstinacy is the better policy)**.** With obstinate policy $\mathcal{O}$, naïve policy $\mathcal{N}$, $1/\rho$ the average arrival rate of information for 3 propositions, $1/\mu$ the expected processing rate, and working memory m:
For m from 1 to 5, $\mathcal{O} >_{30} \mathcal{N}$ whenever $2\rho > \mu$. For $m = 1$ or $m = 2$, $\mathcal{O} >_{30} \mathcal{N}$ for up to $10\rho = \mu$. For $m = 4$, $\mathcal{O} >_{30} \mathcal{N}$ up to $4\rho = \mu$...

‡These assumptions fix the distributions X_i and Y_μ down to their mean occurrence rates.

What we learn from the results

4. In the short to medium term the Obstinate Policy is preferable to the naive one under modest assumptions.
5. Stochastic bottlenecking: Obstinacy will be useful even if we employ other cognitive cost-cutting measures — like heuristic, quick-and-dirty, reasoning.

VI Applications

VI.1 Clarifying the belief-as-plan view

On this view, belief in p:

6. is a paradigmatically stable *epistemic* coordination point or an epistemic quasi-fixed-point; and
7. trades off the epistemic flexibility of being able to constantly fine-tune our epistemic stance towards p by *resisting reconsideration* on the issue in order make up epistemic gains on other issues.

...sound familiar?

Questions that one might have are why this is so, to what extent, and in what circumstances. The *credal* resistance to reconsideration involved in being obstinate can ground an answer.

VI.2 Cognitive "biases"?

The model may help to explain why certain of our information processing practices, termed 'biases' in the cognitive science literature are beneficial.

8. 'Primacy' effects (Peterson and DuCharme 1967)
9. Persistence effects (Ross et al. 1975).

VI.3 Avoiding (some) "demandingness worries"

The outlined policy helps avoid Harman (1986, ch. 3) and Holton's (2013, pp. 2-3, 10-2) worries that credal reasoning requires us to be unrealistically epistemically responsive.