This paper (which adds Ben Eva as its main author) is a sequel to Shear & Fitelson [14], which explored the relationship between probabilistic and qualitative belief *revision* operators.

We will aim to do three main things in this talk:

- Review the results from Shear & Fitelson [14] regarding probabilistic *vs* qualitative belief revision operators.
- Explain the distinction between belief revision *vs* update.
- Compare and contrast (in a style similar to that of [14]) probabilistic *vs* qualitative belief update operators.

We begin with some formal background for belief *revision*.

Our agents will possess *both* qualitative (yes/no) belief sets, $\mathbf{B}$, *and* numerical confidence/credence functions, $b(\cdot)$.

On the *belief* side, our agents entertain (classical, possible worlds) propositions on some finite *agenda* $\mathcal{A}$.

(1) $\mathbf{B}$ is the set of propositions in $\mathcal{A}$ believed by our agent.
   - *Note:* when $p \in \mathbf{B}$, we write $\mathsf{B}(p)$.

(2) Given a *prior* belief set $\mathbf{B}$, the *revised* belief set $\mathbf{B}'$ is generated by revising the prior by E — *i.e.*, $\mathbf{B}' = \mathbf{B} \star E$.

On the *credence* side:

(3) $b(\cdot)$ is a classical (Kolmogorov) probability function.

(4) Given a *prior* $b(\cdot)$, the *revised* $b'(\cdot)$ is generated *via* *conditionalization* on E — *i.e.*, $b'(\cdot) = b(\cdot \mid E)$.[1]

[1] Our results [14] generalize to "minimum distance" [4] Bayesian revisions satisfying (i) $b'(E) > b(E)$, (ii) $b'(E) \geq t$, and (iii) $b'(X) \geq t \Rightarrow b(E \supset X) \geq t$.

The probabilistic approach implies a (*weak*) *Lockean thesis* [5, 7].

$$\mathsf{B}(p) \quad \textit{only if} \quad b(p) \geq t.$$

A *pure Bayesian* is a *strong Lockean* (both *synchronically & diachronically*). Specifically, they satisfy the following:

Pure Bayesian Revision (PBR). When one revises one's beliefs, one's posterior belief set $\mathbf{B}'$ *is Lockean*:

$$\mathbf{B}' = \mathbf{B} \circledast E \stackrel{\text{def}}{=} \{p \mid b(p \mid E) \geq t\}$$

It is well known [2, 8] that *strong* (*iff*) Bayesianism/Lockeanism leads to belief sets (*e.g.*, lotteries) that are *not deductively cogent*.

Cogency. An agent's belief set $\mathbf{B}$ is *cogent* iff it is (a) deductively *consistent* and (b) *closed* under logic.

Shear & Fitelson discussed both **PBR** and **CBR** for Lockean agents (*i.e.*, where **CBR** = **PBR** + **Cogency** [11]).

AGM is the orthodox, *qualitative* account of belief revision.

AGM can be understood as embodying a *principle of* ***conservativity*** (*aka., informational economy/minimal mutilation*).

Conservativity. When an agent with a prior belief set $\mathbf{B}$ learns (exactly) E, she should revise to a posterior belief set $\mathbf{B}'$ that:

(1) *includes* E,

(2) is *cogent*, and

(3) constitutes a *minimal change*[2] to $\mathbf{B}$.

The core formal properties shared by all conservative/AGM belief revision operators can be neatly axiomatized.

[2] "Minimal change" can be explicated (in *geodesic* terms) *via* a wide variety of measures of distance between prior and posterior belief sets [9, 3, 6].

Basic Gärdenfors postulates (AGM axioms):

(*1)	$\mathbf{B} * E = \mathsf{Cn}(\mathbf{B} * E)$	**Closure**
(*2)	$E \in \mathbf{B} * E$	**Success**
(*3)	$\mathbf{B} * E \subseteq \mathsf{Cn}(\mathbf{B} \cup \{E\})$	**Inclusion**
(*4)	If E is consistent with $\mathbf{B}$, then $\mathbf{B} * E \supseteq \mathbf{B}$	**Preservation**
(*5)	If E is consistent, then $\mathbf{B} * E$ is consistent	**Consistency**
(*6)	If $\vDash E_1 \equiv E_2$, then $\mathbf{B} * E_1 = \mathbf{B} * E_2$	**Extensionality**

[Strictly speaking, Gärdenfors' (AGM4) was **Vacuity** and not **Preservation**.

(*4) If E is consistent with $\mathbf{B}$, then $\mathbf{B} \star E \supseteq \mathsf{Cn}(\mathbf{B} \cup \{E\})$ **Vacuity**

But, given the other postulates, these two axioms are equivalent. And, it makes our presentation more elegant/continuous to use **Preservation** here, because it implies both of the weaker forms of Preservation above.]

Shear & Fitelson [14] focused on *preservation*, which is the most interesting source of disagreement between Bayesian and AGM revision. We looked at *three grades* of preservation:

***Very* Weak Preservation.** If $E, X \in \mathbf{B}$, then $\neg X \notin \mathbf{B} \star E$.

Weak Preservation. If $E, X \in \mathbf{B}$, then $X \in \mathbf{B} \star E$.

Preservation. If E is consistent with $\mathbf{B}$, then $\mathbf{B} \star E \supseteq \mathbf{B}$.

Our analysis revealed the following four facts:

(i) **PBR** satisfies ***Very* Weak Preservation**, *if* $t \in (\phi^{-1}, 1]$.

(ii) **PBR** *need not* satisfy **Weak Preservation** — *unless* $t = 1$.

(iii) **CBR** *must* satisfy **Weak Preservation** (for *all* t).

(iv) **CBR** *need not* satisfy **Preservation** — *unless* $t \in [1/2, \phi^{-1})$.

Philippa knows that one of two items (a book and a magazine) is on the table and the other is on the floor (in the next room). Suppose Philippa instructs a robot to enter the next room and make sure that the book is on the floor. The robot will approach the table and if the book is on the table the robot will place it on the floor; otherwise it will do nothing. [16]

Assuming Philippa is cogent, her initial belief set will be:

$$\mathbf{B} = \mathsf{Cn}(\{(B \,\&\, \neg M) \vee (\neg B \,\&\, M)\})$$

where B := 'book on table' and M := 'magazine on table'.

Suppose Philippa learns (after the robot returns) that the book is on the floor ($\neg B$). What should her new belief set look like?

AGM revision entails that $\mathbf{B} * \neg B = \mathsf{Cn}(\{\neg B \,\&\, M\})$.

It seems wrong for Philippa to come to believe M. She knows the robot didn't move the magazine (and it didn't move on its own).

This is not just a problem for AGM revision. It is also a problem for Bayesian revision. Suppose Philippa's initial credences are:

$$b(B \wedge M) = 0, b(B \wedge \neg M) = \frac{1}{2}, b(\neg B \wedge M) = \frac{1}{2}, b(\neg B \wedge \neg M) = 0$$

If she revises her credences by $\neg B$ (*i.e.*, if she *conditionalizes* b on $\neg B$), then she will become certain that M is true, even though she knows that the robot left the magazine untouched.

The best known solution to this problem from the literature is given by Lewis's [12] *imaging* rule for updating probabilities.

Both conditioning and imaging can be understood as "minimal changes" to b, but in different senses of "minimal."

Conditioning on E "preserves the profile of probability ratios, equalities & inequalities among sentences that imply E." [12]

Imaging on E, on the other hand, "involves no gratuitous movement of probability from worlds to dissimilar worlds." [12]

Lewis makes this idea precise by employing Stalnaker's [15] idea of a *selection function* $\sigma(w, X)$, which selects/delivers the closest (or most similar) X-world to a possible world w.

Selection functions are assumed to have 2 basic properties [15].

Centering. If $w \vDash X$, then $\sigma(w, X) = w$.

Uniformity. If $\sigma(w, X) \vDash Y$ and $\sigma(w, Y) \vDash X$, then $\sigma(w, X) = \sigma(w, Y)$.

If we think of possible worlds as binary (**T**/**F**) vectors, then we can define a *Carnapian* selection function $\sigma^{\mathfrak{c}}(w, X)$ as the X-world which *minimizes Hamming distance* from w.[3]

Using σ, we can make Lewis's informal idea that imaging involves "moving probability to similar worlds" precise.

[3] If there is a tie, we can pick the top such w in the truth-table ordering.

Definition

The result $b_E(w)$ of *imaging* a probability function $b(w)$ on E is defined as follows, where $\overline{\mathbf{E}}_w$ is the set of $\neg E$-worlds having w as their closest E-world, *i.e.*, $\overline{\mathbf{E}}_w \overset{\text{def}}{=} \{w' \mid w' \vDash \neg E \,\&\, w = \sigma(w', E)\}$.

$$b_E(w) \overset{\text{def}}{=} \begin{cases} b(w) + \sum\limits_{w' \in \overline{\mathbf{E}}_w} b(w') & \text{if } w \text{ is an } E\text{-world} \\ 0 & \text{if } w \text{ is a } \neg E\text{-world.} \end{cases}$$

Simple example: suppose we have two atomic sentences $\{P, E\}$. Then, we can compare $b(w \mid E)$ and the *Carnapian* $b^{\mathfrak{c}}_E(w)$.

P	E	$b(w)$	$b(w \mid E)$	$b^{\mathfrak{c}}_E(w)$
T	**T**	a	$\frac{a}{a+c}$	$a+b$
T	**F**	b	0	0
F	**T**	c	$\frac{c}{a+c}$	$c+d$
F	**F**	d	0	0

Conditioning *vs* Carnapian imaging — on our example above:

B	M	$b(w)$	$b(w \mid \neg B)$	$b^{\mathfrak{c}}_{\neg B}(w)$
T	**T**	0	0	0
T	**F**	1/2	0	0
F	**T**	1/2	1	1/2
F	**F**	0	0	1/2

Imaging yields a more intuitive credal change than conditioning.

On the qualitative side, a "minimal change" semantics for belief update was developed by Katsuno & Mendelzon [10].

Like Lewis, they move away from distance/similarity between *doxastic states*, and work with distance between *worlds* instead.

They associate with each world w a preorder $\leq_w$, where "$w_1 \leq_w w_2$" means "w_1 is at least as similar to w as w_2 is."

A preorder $\leq_w$ is said to be *faithful* if it satisfies both:

1. For all w, w': $w \leq_w w'$ and if $w \neq w'$, then $w <_w w'$. That is, w is *the minimal element of* $\leq_w$.
2. For every consistent p and every world w, the set $[p]$ has at least one minimal element with respect to $\leq_w$. Where $[p]$ is the set of worlds compatible with p (more generally, $[\mathbf{B}]$ is the set of worlds compatible with all of the members of $\mathbf{B}$).

Then, KM define *update* of a belief set $\mathbf{B}$ by a sentence E (in language $\mathscr{L}$), based on a faithful assignment $\leq_w$, as follows:

$$\mathbf{B} \circ E \overset{\text{def}}{=} \begin{cases} \bigcup\limits_{w \in [\mathbf{B}]} \min([E], \leq_w) & \text{if } [\mathbf{B}] \neq \varnothing \text{ and } [E] \neq \varnothing \\ \mathscr{L} & x`pt\text{otherwise} \end{cases}$$

$\min([p], \leq_w)$ is the set of minimal elements in $[p]$ with respect to $\leq_w$, *i.e.*, $\min([p], \leq_w) \overset{\text{def}}{=} \{z \in [p] \mid \nexists z' \in [p] \text{ s.t. } z' <_w z\}$.

KM show that the following axioms are sound and complete with respect to this “minimal change” update semantics [10, 13].

($\circ 1$) $\mathbf{B} \circ E$ is closed under logical entailment

($\circ 2$) $E \in \mathbf{B} \circ E$

($\circ 3$) If $E \in \mathbf{B}$ then $\mathbf{B} = \mathbf{B} \circ E$

($\circ 4$) If $\mathbf{B}$ and E are individually consistent then $\mathbf{B} \circ E$ is consistent

($\circ 5$) If $\vDash E_1 \equiv E_2$, then $\mathbf{B} \circ E_1 = \mathbf{B} \circ E_2$

($\circ 6$) $\mathbf{B} \circ (E_1 \wedge E_2) \subseteq \mathsf{Cn}\left((\mathbf{B} \circ E_1) \cup \{E_2\}\right)$

($\circ 7$) If $E_1 \in \mathbf{B} \circ E_2$ and $E_2 \in \mathbf{B} \circ E_1$, then $\mathbf{B} \circ E_1 = \mathbf{B} \circ E_2$

($\circ 8$) If $\mathbf{B}$ is complete, then $\mathbf{B} \circ (E_1 \vee E_2) \subseteq \mathsf{Cn}((\mathbf{B} \circ E_1) \cup (\mathbf{B} \circ E_2))$

($\circ 9$) $\mathbf{B} \circ E = \bigcap_{K \in \mathscr{K}_{\mathbf{B}}} K \circ E$

Where $\mathscr{K}_{\mathbf{B}}$ is the set of complete, cogent belief sets whose strongest members are compatible with all members of **B**.

Pure Bayesian belief update is defined as follows:

$$B \diamond E \overset{\text{def}}{=} \{p \mid b_E(p) \geq t\}, \text{ for some } t \in (1/2, 1]$$

All told, we’ve seen *six* kinds of belief change.

Pure Bayesian revision (**PBR**): $B \circledast E \overset{\text{def}}{=} \{p \mid b(p \mid E) \geq t\}$.

Cogent Bayesian revision (**CBR**): **PBR** + **Cogency**.

Pure Bayesian update (**PBU**): $B \diamond E \overset{\text{def}}{=} \{p \mid b_E(p) \geq t\}$.

Cogent Bayesian update (**CBU**): **PBU** + **Cogency**.

AGM: $B * E$, in accordance with AGM axioms/“minimal change.”

KM: $B \circ E$, in accordance with KM axioms/“minimal change.”

We have now characterized the relationships between Bayesian revisions/updates and **AGM/KM.** I’ll conclude with a summary.

	AGM						**KM**					
	$*1$	$*2$	$*3$	**VWP** $\mid$ **WP** $\mid$ $*4$	$*5$	$*6$	$\circ 3$	$\circ 4$	$\circ 6$	$\circ 7$	$\circ 8$	$\circ 9$
PBR	×	✓	✓	$t > \phi^{-1} \mid \times \mid \times$	×	✓	×	×	✓	×	×	?
CBR	✓	✓	✓	$\checkmark \mid \times \mid t < \phi^{-1}$	✓	✓	✓	✓	✓	×	$t < \phi^{-1}$	?
PBU	×	✓	✓	$t > 2/3 \mid \times \mid \times$	×	✓	×	×	✓	×	✓	?
CBU	✓	✓	✓	$\checkmark \mid \checkmark \mid \times$	✓	✓	✓	✓	✓	✓	✓	✓

Note: $\circ 1$, $\circ 2$, and $\circ 5$ are omitted, since $\circ 1 = *1$, $\circ 2 = *2$, and $\circ 5 = *6$.

[1] C. Alchourron, P. Gärdenfors, and D. Makinson. *On the logic of theory change: Partial meet contraction and revision functions*, 1985.

[2] D. Christensen, *Putting Logic in its Place*, 2007.

[3] M. Deza and E. Deza, *Encyclopedia of Distances*, 2009.

[4] P. Diaconis and S. Zabell, *Updating Subjective Probability*, 1982.

[5] K. Dorst, *Lockeans Maximize Expected Accuracy*, 2015.

[6] C. Duddy and A. Piggins, *A measure of distance between judgment sets*, 2012.

[7] K. Easwaran, *Dr. Truthlove or: How I Learned to Stop Worrying and Love Bayesian Probability*, 2015.

[8] R. Foley, *Working Without a Net*, 1992.

[9] K. Georgatos, *Geodesic Revision*, 2008.

[10] H. Katsuno and A. Mendelzon, *On the difference between updating a knowledge base and revising it*, 1992.

[11] H. Leitgeb. *Stability Theory of Belief.* Oxford University Press, 2016.

[12] D. Lewis, *Probabilities of Conditionals and Conditional Probabilities*, 1976.

[13] P. Pappas, *Belief Revision*, 2008.

[14] T. Shear & B. Fitelson, *Two Approaches to Belief Revision*, 2018.

[15] R. Stalnaker, *A Theory of Conditionals*, 1968.

[16] M. Winslett, *Reasoning about action using a possible models approach*, 1988.