Four Approaches to Supposition

Benjamin Eva[1] Ted Shear[2] Branden Fitelson[3]

[1]University of Konstanz
benjamin.eva@uni-konstanz.de

[2]University of Queensland
me@tedshear.org

[3]Northeastern University
branden@fitelson.org

PRICAI 2019
Knowledge Representation Conventicle
Foundations of Belief Change

August 27, 2019

As Joyce [3, p. 182] tells us:

> *"In general, a supposition is a form of provisional belief revision in which a person accepts some proposition as true and makes the minimum changes in her other opinions needed to accommodate this modification."*

In categorical contexts, we rely on our ordinary current opinions to assess propositions, while suppositions trigger shifts in the modal base.

We contrast canonical examples of four types of suppositional theories.

Consider the following minimal pair differing only in grammatical mood:

1. Suppose that Oswald *didn't* shoot Kennedy...
   - I may then infer that someone else did.
2. Suppose that Oswald *hadn't* shot Kennedy...
   - I may then infer that Kennedy would have left Dallas unharmed.

This suggests a semantic difference between two modes of supposition.

Indicative/'epistemic' suppositions trigger revisions that align those that would follow from learning that the supposition is true.

Subjunctive/'ontic' suppositions elicit the agent's counterfactual opinions about how things would be if the supposition were suddenly made true through a 'local miracle' or 'ideal intervention'.

Indeed, different belief change operations are needed to characterise the shifts triggered by indicative and subjunctive suppositions.

Theories of supposition can be partitioned along another dimension.

Qualitative approaches give binary *acceptability* judgements based in part on agents' qualitative/outright/categorical beliefs.

- Conditional logics, belief revision theories, and non-monotonic logics have all been proposed as qualitative theories of supposition.

Quantitative approaches give graded *degree of acceptability* judgements based in part on agents' numerical credences/subjective probabilities.

- Accounts of quantitative theories have largely been developed and explored in the context of decision theory.

In sum, this leaves four species of suppositional theories:

Qualitative-Indicative (AGM Revision)	Quantitative-Indicative (Conditionalization)
Qualitative-Subjunctive (KM Update)	Quantitative-Subjunctive (Imaging)

On the indicative side: we contrast AGM revision with conditionalization.

On the subjunctive side: we contrast KM update with the imaging rule.

The intuitive connections between these accounts are often drawn—*e.g.* in their seminal paper, Katsuno and Mendelzon [4, p.184] write:

> *"We can regard imaging as a probabilistic version of update and conditionalization as a probabilistic version of revision."*

Before proceeding, some formalism:

- Given some (finite) propositional language $\mathcal{L}$, let $\mathcal{A}$ be an algebra of classical possible-world propositions expressible in $\mathcal{L}$.
- The agent's beliefs are given by $\mathbf{B} \subseteq \mathcal{L}$ (*not* subject to any special requirements unless otherwise specified).
- The agent's credences are represented by a classical probability function $b : \mathcal{A} \mapsto [0, 1]$.
- Qualitative theories are given of operations on agents' beliefs and formulas to *acceptability sets*—*e.g.* $\mathbf{B} \circ S$ contains everything $\circ$ says are acceptable under supposition S for an agent with beliefs $\mathbf{B}$.
- Quantitative theories generate *acceptability functions* from agents' credences together with a supposition—*e.g.* $b^S(\cdot)$ gives numerical degrees of acceptability under S for an agent with credences $b(\cdot)$.

Our approach differs from previous explorations of the connections between probabilistic and qualitative revision such as [6, 7, 9, 1].

Instead, we rely on the simplistic threshold-based account from the *Lockean Thesis* as a bridge principle allowing us to construct qualitative judgments that cohere with underlying quantitative ones.

$$\text{For some } t \in [1/2, 1] : X \in \mathbf{B} \Leftrightarrow b(X) > t \qquad (\text{LT}^t)$$

The Lockean Thesis is often discussed (esp. in the context of the Lottery paradox) as a *synchronic coherence requirement* on credence & belief.

Generalising LT^t to suppositional judgments lets us define operators driven by quantitative accounts as follows:

$$\mathbf{B} \circ^t S := \{X : b^S(X) > t\}$$

In general, our approach is in tension with the *knowledge level* representation presupposed by qualitative theories.

This presupposition corresponds to requiring that **B** is cogent.

Cogency: A set of belief **B** is *cogent* just in case:

1. **B** is consistent, *i.e.* $\bot \notin \text{Cn}(\mathbf{B})$, and
2. **B** is deductively closed, *i.e.* $\mathbf{B} = \text{Cn}(\mathbf{B})$.

We set this issue aside as largely orthogonal to our more general interests in the relative behaviour of Lockean and qualitative theories.

So, while we still consider how things fare without **Cogency**, our most interesting results will involve cogent Lockean agents.

We start with the quantitive theory given by Bayesian conditionalization:

$$b^S(X) := b(X \mid S) = \frac{b(X \wedge S)}{b(S)}$$

Applying the scheme from earlier lets us define the $\divideontimes$ operator:

$$\mathbf{B} \divideontimes^t S := \{X : b(X \mid S) > t\}$$

Some conventions:

- we write $\divideontimes$ when the threshold is irrelevant, and
- we let $\divideontimes^{[m,n]}$ to denote the class of $\divideontimes^t$ operators such that $m \leq t \leq n$.

Next, we consider how $\divideontimes$ holds up under the Gärdenfors postulates.

The Gärdenfors postulates below axiomatise AGM's revision operator:

(*1) $\mathbf{B} * S = \text{Cn}(\mathbf{B} * S)$ (*Closure*)
(*2) $S \in \mathbf{B} * S$ (*Success*)
(*3) $\mathbf{B} * S \subseteq \mathbf{B} + S$ (*Inclusion*)
(*4) If $\neg S \notin \text{Cn}(\mathbf{B})$, then $\mathbf{B} \subseteq \mathbf{B} * S$ (*Preservation*)
(*5) If S is consistent, then $\mathbf{B} * S$ is consistent (*Consistency*)
(*6) If $S \Leftrightarrow S'$, then $\mathbf{B} * S = \mathbf{B} * S'$ (*Extensionality*)

In earlier work [8], we established that *Preservation* is at the center of the difference between AGM and Lockean revision as seen below.

	*1	*2	*3	*4	*5	*6
$\divideontimes$		✓	✓			✓
$\divideontimes$ + **Cogency**	✓	✓	✓		✓	✓

But, things are a bit more subtle (and surprising) as we can see by inspecting two weakened versions of *Preservation*:

(*4^v) If $E, X \in \mathbf{B}$, then $\neg X \notin \mathbf{B} * E$ (*Very Weak Preservation*)
(*4^w) If $E \in \mathbf{B}$, then $\mathbf{B} \subseteq \mathbf{B} * E$ (*Weak Preservation*)

We find a notable bound on t at the inverse golden ratio ($\Phi \approx 0.618$).

	*4^v	*4^w	*4
$\divideontimes^{[\Phi,1]}$	✓		
$\divideontimes$ + **Cogency**	✓	✓	
$\divideontimes^{[1/2,\Phi)}$ + **Cogency**	✓	✓	✓

This is largely since Φ is the unique positive value such that $x^2 = 1 - x$.

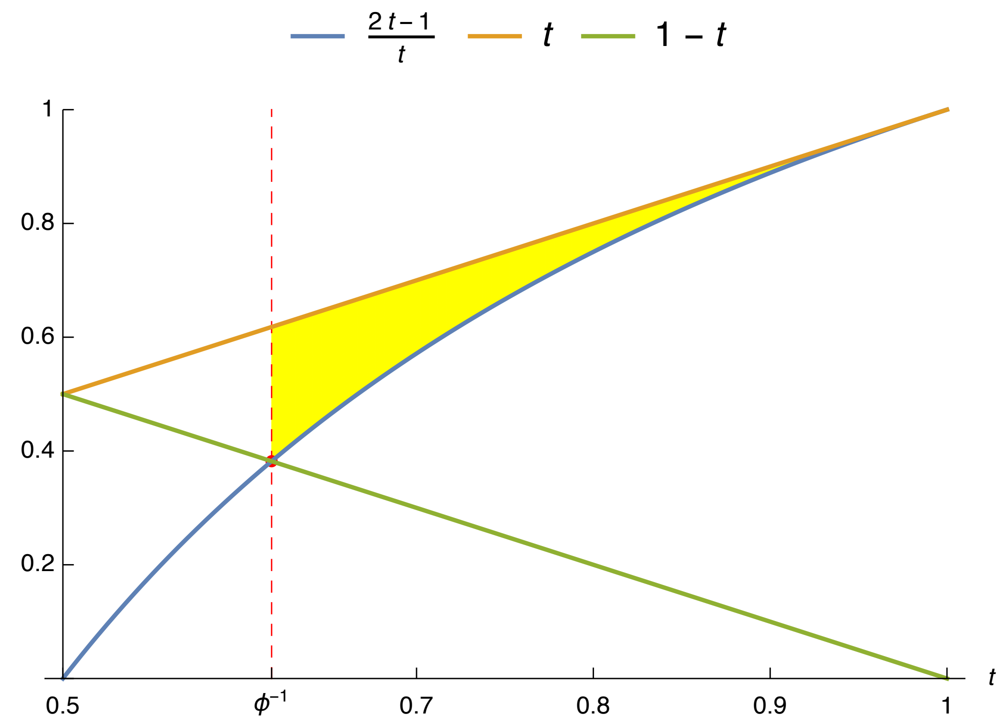

Figure: Graphical explanation of why $\divideontimes^{[\Phi,1]}$ satisfies (*4^v)

For subjunctives, we consider the quantitative account given by imaging.

As Lewis [5, p. 311] described it:

> *"Imaging [. . .] gives a minimal revision in this sense: unlike all other revisions [. . .], it involves no gratuitous movement of probability from worlds to dissimilar worlds."*

Imaging relies on a *selection function* $\sigma : W \times \mathcal{L} \mapsto W$ subject to the following conditions familiar from the Stalnaker conditional.

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

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

So, $\sigma(w, X)$ is the *unique* closest world to w satisfying X.

Formally, (general) imaging b on X, denoted $b_X(\cdot)$, is defined:

$$b_S(w) := \begin{cases} b(w) + \sum\limits_{\substack{u \in [\![\neg S]\!] \\ w = \sigma(u,S)}} b(u) & \text{if } w \in [\![S]\!] \\ 0 & \text{otherwise} \end{cases}$$

This shifts the probability mass from to each w to the nearest X-world.

Applying our scheme from earlier we define the $\blacklozenge$ operator:

$$\mathbf{B} \blacklozenge S := \{X : b_S(X) \geq t\}$$

Now, we consider $\blacklozenge$ in terms of the KM postulates for update.

Consider the syntactic formulation of the KM postulates [4]:

($\diamond$0) $\mathbf{B} \diamond S = \text{Cn}(\mathbf{B} \diamond S)$

($\diamond$1) $S \in \mathbf{B} \diamond S$

($\diamond$2) If $S \in \mathbf{B}$ then $\mathbf{B} = \mathbf{B} \diamond S$

($\diamond$3) If $\mathbf{B}$ and S are consistent, then $\mathbf{B} \diamond S$ is consistent

($\diamond$4) If $\vdash S_1 \equiv S_2$, then $\mathbf{B} \diamond S_1 = \mathbf{B} \diamond S_2$

($\diamond$5) $\mathbf{B} \diamond (S_1 \wedge S_2) \subseteq (\mathbf{B} \diamond S_1) + S_2$

($\diamond$6) If $S_1 \in \mathbf{B} \diamond S_2$ and $S_2 \in \mathbf{B} \diamond S_1$, then $\mathbf{B} \diamond S_1 = \mathbf{B} \diamond S_2$

($\diamond$7) If $\mathbf{B}$ is complete, then $\mathbf{B} \diamond (S_1 \vee S_2) \subseteq (\mathbf{B} \diamond S_1) + (\mathbf{B} \diamond S_2)$

($\diamond$8) $(\mathbf{B}_1 \cap \mathbf{B}_2) \diamond S = (\mathbf{B}_1 \diamond S) \cap (\mathbf{B}_2 \diamond S)$

Like with $\circledast$ and (*1)–(*6), there's not full correspondence.

	$\diamond$0	$\diamond$1	$\diamond$2	$\diamond$3	$\diamond$4	$\diamond$5	$\diamond$6	$\diamond$7	$\diamond$8
$\blacklozenge$		✓			✓	✓		✓	✓

Note: ($\diamond$8) is analogous to the fact established by Gardenförs [2, p. 113] that a probability update is an instance of imaging just in case it preserves mixtures, *i.e.*

$$[\alpha P(X) + (1 - \alpha)P'(X)]_S = \alpha P_S(X) + (1 - \alpha)P'_S(X)$$

We strengthen **Cogency** to include closure under the Stalnaker conditional (→), where $w \models X \rightarrow Y \Leftrightarrow \sigma(w, X) \models Y$.

→-Cogency: A set of belief **B** is →-*cogent* just in case:

1. **B** is consistent, *i.e.* $\perp \notin$ Cn(**B**),
2. **B** is deductively closed, *i.e.* **B** = Cn(**B**), and
3. **B** is closed under →, *i.e.* $X, X \rightarrow Y \in$ **B** implies $Y \in$ **B**.

With this in place, we get a general satisfaction of the KM postulates.

	◇0	◇1	◇2	◇3	◇4	◇5	◇6	◇7	◇8
◆ + →-**Cogency**	✓	✓	✓	✓	✓	✓	✓	✓	✓

We've now made two pairwise comparisons:

Indicative Qualitative (AGM Revision)	Indicative Quantitative (Conditionalization)
Subjunctive Qualitative (KM Update)	Subjunctive Quantitative (Imaging)

We found an even tighter parallel between imaging and KM update than between conditionalization and revision as no special threshold was needed.

To conclude, we consider the remaining diagonal comparisons.

For ◆ and the AGM postulates, we find that there is no threshold for which (*4) is satisfied (even under the assumption of →-**Cogency**).

	*1	*2	*3	*4	*5	*6
◆ + →-**Cogency**	✓	✓	✓		✓	✓
⋇$^{[1/2,\Phi]}$ + **Cogency**	✓	✓	✓	✓	✓	✓

The situation is also interesting for ⋇ and the KM postulates:

	◇0	◇1	◇2	◇3	◇4	◇5	◇6	◇7	◇8
⋇		✓			✓	✓			
⋇ + **Cogency**	✓	✓	✓	✓	✓	✓	✓		
⋇$^{[1/2,\Phi]}$ + **Cogency**	✓	✓	✓	✓	✓	✓	✓	✓	

Here, (◇7) is satisfied only under when **Cogency** and $t \in [1/2, \Phi]$ are assumed, while there is no threshold such that ⋇ will satisfy (◇8).

We gave a Lockean vindication to the widely claimed parallel between revision/conditionalization and update/imaging.

- ⋇ satisfies the AGM postulates when **Cogency** is assumed and t is restricted to $[1/2, \Phi]$.
- ◆ satisfies the KM postulates when **Cogency** is extended to include closer under the Stalnaker conditional.

In future work, we plan to explore:

- Lockean operators based on other types of imaging aside from general imaging,
- Lockean operators from a semantic perspective,
- explicit suppositional reasoning and choice for Lockean agents,and
- other varieties of Lockean operators.

[1] K. Chhogyal, A. Nayak, R. Schwitter, and A. Sattar. Probabilistic belief revision via imaging. In D. Pham and S. B. Park, editors, *PRICAI 2014, LNAI 8862*, pages 694–707, 2014.

[2] P. Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. Number 13 in Studies in Logic. MIT Press, Cambridge, 1988.

[3] J. M. Joyce. *The Foundations of Causal Decision Theory*. Cambridge Studies in Probability, Induction, and Decision Theory. Cambridge University Press, Cambridge, 1999.

[4] H. Katsuno and A. O. Mendelzon. On the difference between updating a knowledge base and revising it. In P. Gärdenfors, editor, *Belief Revision*, pages 183–203. Cambridge University Press, 1992.

[5] D. Lewis. Probabilities of conditionals and conditional probabilities. *The Philosophical Review*, 85(3):297–315, 1976.

[6] R. Ramachandran, A. C. Nayak, and M. A. Orgun. Belief erasure using partial imaging. In J. Li, editor, *AI 2010: Advances in Artificial Intelligence*, pages 52–61, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg.

[7] G. Rens, T. Meyer, and G. Casini. On revision of partially specified convex probabilistic belief bases. In *Proceedings of the Twenty-second European Conference on Artificial Intelligence*, ECAI'16, pages 921–929, Amsterdam, 2016.

[8] T. Shear and B. Fitelson. Two approaches to belief revision. *Erkenntnis*, 84:487–518, 2019.

[9] Z. Zhuang, J. Delgrande, A. Nayak, and A. Sattar. A unifying framework for probabilistic belief revision. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI-17*, pages 1370–1376, 2017.